CPU Heads To The Crossroads Guitar Festival (page 94) ■ Intel 775 Platform Reviewed (page 20)

A Sandhills Publication

CPU®

COMPUTER POWER USER

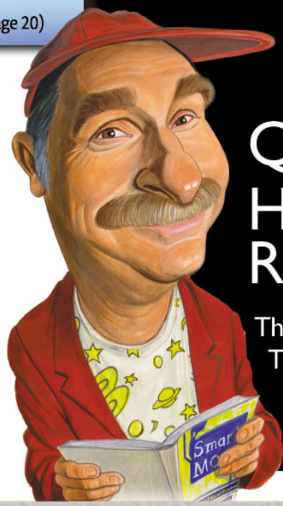

Q&A With Howard Rheingold

The "First Citizen Of The Internet" Shares His Thoughts (page 108)

Alienware ALX X2
Two PCI-E Slots Are Nicer Than One (page 46)

Broadband Phones
We Roundup The Services (page 71)

Macromedia's Flash
Struts Its Stuff (page 80)

THE NEXT BIG TECH BREAKTHROUGHS

page 52

Artificial Intelligence
Biometrics
Biotechnology
Cooling
Grid Computing
MEMS & Nanobots
OLED & Flexible Displays
Optical Computing
Quantum Computing
Robotics

Check Out This Cammer V8 Modded PC
On Page 36

August 2004 • Vol. 4 Iss. 8

$5.99 U.S.
$7.99 Canada

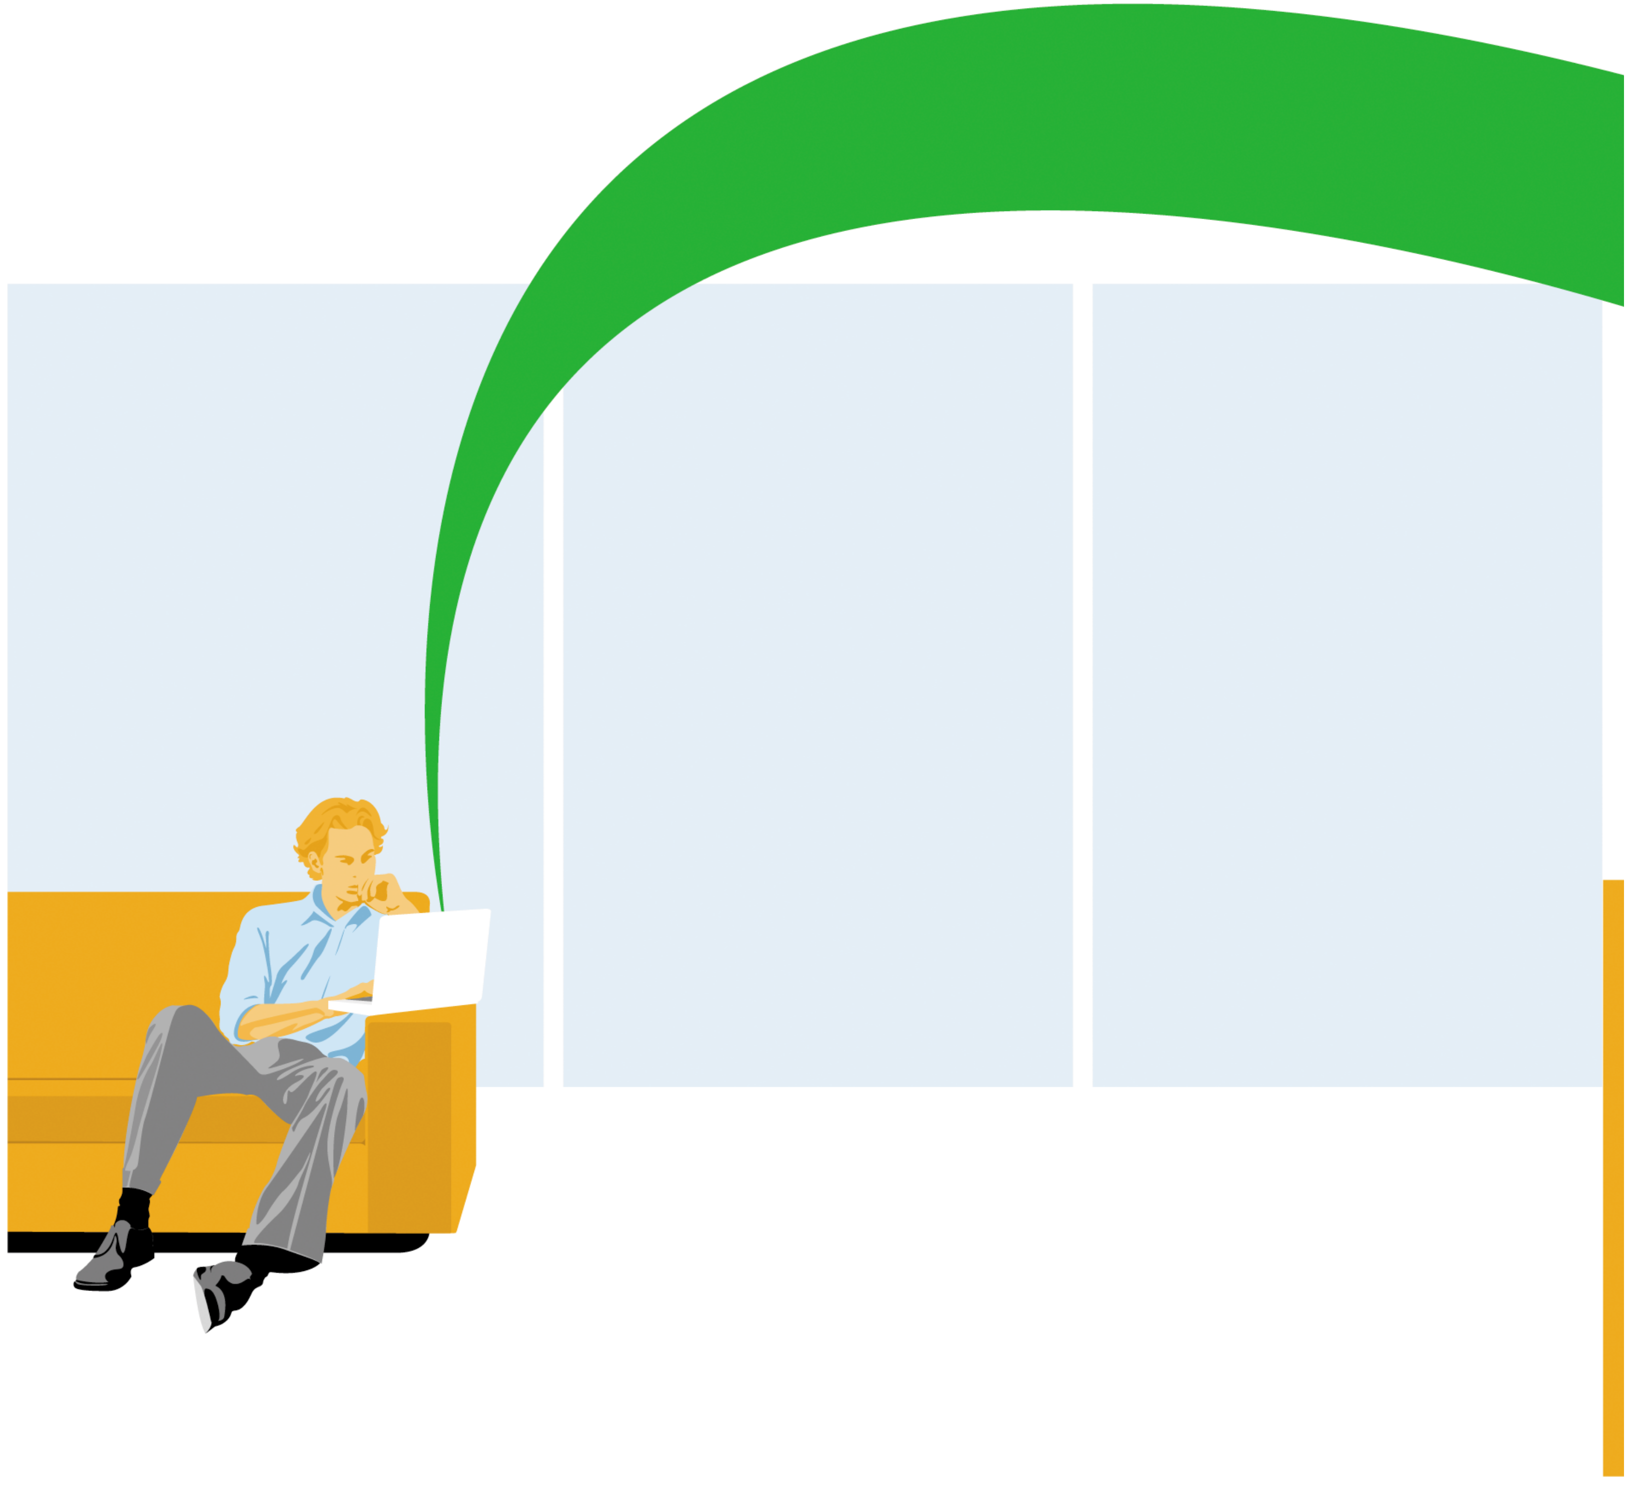

GoToMyPC®

Access Your PC from Anywhere™

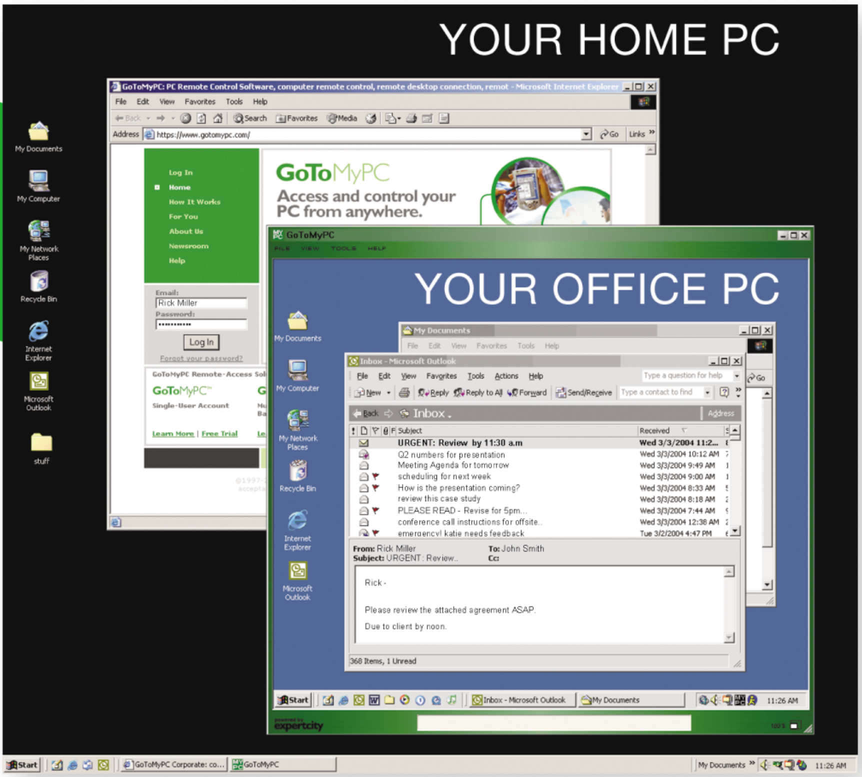

YOUR HOME PC
GoToMyPC
Access and control your PC from anywhere.
YOUR OFFICE PC
Inbox - Microsoft Outlook
My Documents

URGENT: Review by 11:30 a.m — Wed 3/3/2004 11:2...
Q2 numbers for presentation — Wed 3/3/2004 10:12 AM
Meeting Agenda for tomorrow — Wed 3/3/2004 9:49 AM
scheduling for next week — Wed 3/3/2004 9:00 AM
How is the presentation coming? — Wed 3/3/2004 8:33 AM
review this case study — Wed 3/3/2004 8:18 AM
PLEASE READ - Revise for 5pm... — Wed 3/3/2004 7:44 AM
conference call instructions for offsite.. — Wed 3/3/2004 12:38 AM
emergency! katie needs feedback — Tue 3/2/2004 4:47 PM

From: Rick Miller
To: John Smith
Subject: URGENT: Review..

Rick -

Please review the attached agreement ASAP.

Due to client by noon.

368 Items, 1 Unread

- Full Web-based remote access and control
- No pre-installed client software needed
- Automatic 2-minute one-time setup
- Works automatically with firewalls
- 128-bit end-to-end AES encryption
- Real-time desktop collaboration
- FREE tech support & upgrades for life

Try It Free

gotomypc.com/cpu3

888.666.7490

CPU
COMPUTER POWER USER®

August 2004 • Vol 4 Issue 8

Spotlight

52 **Big Technology Breakthroughs**
A Lesson In Computing History

56 **Artificial Intelligence**: Boosting Your Brain With Software Assistants
57 **Biometrics**: Your Body Will (Un)lock Your World
58 **Biotechnology**: Can Technology Succeed Where Mother Nature Has Failed?
60 **Cooling**: Dealing With Faster, More Powerful Processors
61 **Grid Computing**: Not Just Folding@Home Anymore
62 **MEMS & Nanobots**: Ultra-Tiny Machines Will Reshape The World
63 **OLED/Flexible Displays**: Tomorrow's Monitors: Better, Brighter & Bendable
64 **Optical Computing & Switching In Silicon**: Computing Power & Connectivity At Light Speed
65 **Quantum Computing**: Explaining The Inexplicable & Harnessing Its Power
66 **Robotics**: Lifelike Robotic Personal Servants Are No Longer Science Fiction

Frontside

6 **What's Happening**
12 **Digital Economy**
14 **The Saint**
What's In A Qubit?

The Experts

Alex St. John
The Saint
page 14

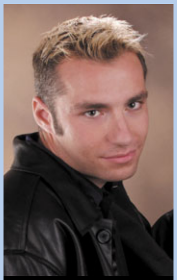
Alex "Sharky" Ross
The Shark Tank
page 27

Joan Wood
Forward Slash
page 87

Anand Lal Shimpi
Anand's Corner
page 26

Rob "CmdrTaco" Malda
The Department Of Stuff
page 86

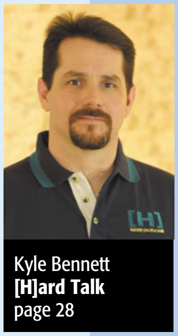
Kyle Bennett
[H]ard Talk
page 28

Mike Magee
Shavings From The Rumour Mill
page 102

Chris Pirillo
Dialogue Box
page 77

Pete Loshin
Open Sauce
page 78

Copyright 2004 by Sandhills Publishing Company. *Computer Power User* is a trademark of Sandhills Publishing Company. All rights reserved. Reproduction of material appearing in *Computer Power User* is strictly prohibited without written permission. Printed in the U.S.A. GST # 123482788RT0001 (ISSN 1536-7568) *CPU Computer Power User* USPS 020-801 is published monthly for $29 per year by Sandhills Publishing Company, 131 West Grand Drive, P.O. Box 82667, Lincoln, NE 68501. Subscriber Services: (800) 424-7900. Periodicals postage paid at Lincoln, NE. POSTMASTER: Send address changes to *Computer Power User*, P.O. Box 82667, Lincoln, NE 68501.

Did you find the hidden CPU logo on our cover? Turn the page for the answer.

Heavy Gear

15 **Dream Hardware**

16 **The Search Continues**
Looking For The Perfect HTPC Case

20 Intel 775 Platform:
915P/G & 925X Chipsets Have Arrived

21 AMD Athlon 64 3800+ & FX-53

22 FIC Condor
Zalman Reserator 1

23 ATI HDTV Wonder
eVGA Personal Cinema FX 5700

24 EPoX EP-8KDA3+ vs. MSI K8N Neo2 Platinum

25 Logitech MX510 Performance Optical Mouse
Cooler Master CM Stacker STC-T01

26 **Anand's Corner**
Intel's New 925X & 915

27 **The Shark Tank**
MXM: Mobile GPU Swap?

28 **[H]ard Talk**
WoOD$+OcK PhoR G33k5

Hard Hat Area

30 PC Modder Tips & Tutorials
31 **Building "Spike"**
The Diary Of A Modder's Creation, Part II

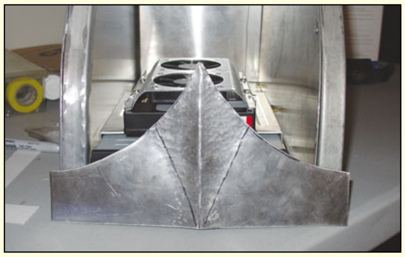

36 **Mad Reader Mod**
This Will Get Your Motor Running

40 **Advanced Q&A Corner**
42 **X-ray Vision**: BIC-TCP
Giving Researchers High-Speed Hopes
46 **White Paper:** Alienware ALX X2
Double Your Pleasure, Double Your Fun, Double Your Video Cards

Loading Zone

68 **The Bleeding Edge Of Software**
Inside The World Of Betas
70 **Up To Speed**
Upgrades That'll Keep You Humming Along
71 **Broadband Phone Services**
The Telephone Meets The Internet
76 MightSoft Audio Editor Pro 1.40
TuneUp Software TuneUp Utilities 2004
77 **Dialogue Box**
Remotely Interested?
78 **Open Sauce**
Going Open Source Is Cheaper

Caught In The Web

80 **Not Just A Flash In The Pan**
Macromedia's Design Tool Hits The Big Time
84 **Coder's Corner: XML**
XML Schema, Part II:
Global & Local Types
86 **The Department Of Stuff**
bizarro.txt
87 **Forward Slash**
Row, Row, Row Your Boat

Digital Living

88 **Road Warrior**
Broadband Over Television, Free Ringtones, Motion-Sensitive Nokia Accessory & More From The Mobile Front
90 **At Your Leisure**
Plug In, Sit Back & Fire Away
94 **Destination Dallas**
CPU Goes To The Crossroads

Tips & Tricks

97 **Software Tips & Projects**
Editing Audio Archives, Part 2
99 **Warm Up To Penguins**
Monitor Your Hard Drive Space

What's Cooking

102 **Shavings From The Rumour Mill**
Will Multicores Save The Game For Intel & AMD?
103 **Technically Speaking**
An Interview With Charlie Boswell, AMD's Director Of Digital Media & Entertainment
106 **Under Development**
A Peek At What's Brewing In The Laboratory

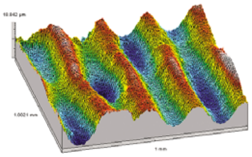

Back Door

108 **Q&A With Howard Rheingold**
The "First Citizen Of The Internet" Shares His Thoughts

Infinite Loops

Strange stats and other oddball items from computing's periphery

81, 98, 100

Greetings From Samitland

Welcome back, folks. The long, hot summer months don't feature much in the way of big trade shows or massive product announcements, but they give us the opportunity to look into topics that we sometimes have trouble finding space for otherwise. On page 52 you'll find our 13-page section on some of the biggest technological breakthroughs around. These are things that will change computing and technology as we know it. Some of them may be considered bordering on sci-fi, but we'd beg to differ. After all, today's technology would have been termed sci-fi less than 100 years ago.

There are some other highlights that I think you'll really find exciting. Check out our more in-depth coverage of what's inside Alienware's upcoming ALX X2 (page 46). Get smart on BIC-TCP, a potentially powerful high-speed protocol that may just change the way researchers use the Internet. We also get our hands dirty with lots of good hardware. How do AMD's Athlon 64 3800+ and FX-53, Intel's 775 platform, Zalman's Reserator 1, and ATI's ATI HDTV Wonder sound? We've got that and more. We've also got some arts and technology coverage this month. For example, take a look at what's being done with Flash on page 80 and read about our trip to the Crossroads Guitar Festival in Dallas. Yeah, we love concerts as much as you, but we went to see how AMD is using its Opterons to rock the casbah (and this historical event). We were impressed, and you can find out why on page 94.

We got a lot of good feedback earlier this year regarding *CPU*'s special issue: *PC Modder*. That's why we now have Son Of PC Modder (officially known as *PC Modder v1.1*). Expect it on newsstands in late July. I know this news will give quite a few of you cause to celebrate, right? Naturally, we'd be quite appreciative if you add a copy of v1.1 to your shopping carts when you're out shopping for the celebratory condiments. And make sure you let us know what you like and don't like so we can continue to improve the mag.

So that's the good news. I also have some bad news. This is the last issue that our very own Jennie Schlueter will be working on. She'll be waving au revoir as she makes her move down to hot, soggy Missouri from hot, humid Nebraska. She's been a valuable member of our team and a good friend, so we're all sad to see her go. Her husband took a new job, and she chose him and her baby over us. We were shocked at her decision, just as you probably are! More seriously: Jennie, we'll miss you. (Tissues are on the house.) Stepping in for Jennie is Dana Montey. She's going to have her hands full, but she's been working on *CPU* as a copy editor for some time now, and I'd like to welcome her aboard in this new capacity.

That's about it for this month. Not like I'm busy or anything, but wife Julie has me doing the type of errands that have her in the running for Most Irritating Spouse Of The Month. (I may have taken the award in a prior month. *Maybe*.) Have a great summer, and see you back here in the September issue. I'm looking forward to it.

Samit G. Choudhuri, Publication Editor, *CPU*

Sandhills Publishing

Customer Service
(For questions about your subscription or to place an order or change an address.)
customer-service@cpumag.com
800-733-3809
FAX: (402) 479-2193
Computer Power User
P.O. Box 82667
Lincoln, NE 68501-2667

Hours
Mon. - Fri.: 7 a.m. to 8 p.m. (CST)
Sat.: 8 a.m. to 4 p.m. (CST)
Online Customer Service & Subscription Center
http://www.cpumag.com

Web Services
(For questions about our Web site.)
webhelp@cpumag.com
(800) 368-8304

Authorization For Reprints
(800) 334-7458
FAX: (402) 479-2104

Editorial Staff
editor@cpumag.com
FAX: (402) 479-2104
131 W. Grand Drive
Lincoln, NE 68521

Subscription Renewals
(800) 382-4552
FAX: (402) 479-2193
http://www.cpumag.com

Advertising Staff
(800) 848-1478
120 W. Harvest Dr.
Lincoln, NE 68521
FAX: (402) 479-2193

The Gang
Editorial Staff: Ronald D. Kobler / Samit Gupta Choudhuri / Jennie Schlueter / Blaine Flamig / Corey Russman / Rod Scher / Christopher Trumble / Calvin Clinchard / Kimberly Fitzke / Katie Dolan / Raejean Brooks / Rebecca Christensen / Tara Weber / Sally Curran / Michael Sweet / Katie Crooks / Shawn Kupfer / Dana Montey / Nate Hoppe / Sheila Allen / Jennifer Hansen / Trista Kunce / Marty Sems / Chad Denton / Nathan Chandler / Kylee Dickey / Josh Gulick / Andrew Leibman / Eric Reed / Vince Cogley **Web staff:** Missy Fletcher / Dorene Krausnick / Nick Ray **Customer Service:** Alisha Lamb / Brandie Humphrey / Becky Rezabek / Lana Matic **Subscription Renewals:** Liz Kohout / Connie Beatty / Matt Bolling / Patrick Kean / Charmaine Vondra / Miden Ebert / Kathy DeCoito / Stephanie Contreras / Nicole Buckendahl / Travis Brock **Art & Design:** Lesa Call / Fred Schneider / Carrie Benes / Ginger Riley / Sonja Warner / Leigh Trompke / Aaron Weston / Aaron Clark / Lori Garris / Jason Codr / Andria Schultz / Erin Rodriguez / Lindsay Anker **Newsstand:** Garth Lienemann / Kelly Richardson / Chris McGreer / Jeff Schnittker **Advertising Sales:** Grant Ossenkop / Cindy Pieper / Brooke Wolzen / Eric Cobb / Emily Getzschman **Marketing:** Mark Peery / Marcy Gunn / Heather Mentgen / Amber Coffin / Jen Clausen

CORRECTIONS

In June's issue (page 106), Komag expects DTR to facilitate platter capacities of 160GB to 240GB or greater, with 1TB per drive, not 4TB.

Gotcha.
Here it is.

newegg.com™

- 8,800+ Customer Testimonials
- A Rating on ANY Merchant Ranking Site
- 98% of Orders Ship within 1 Business Day
- 93,000+ Product Reviews From Customers

FREE Shipping

Best Value, Always!

FREE SHIPPING*
FedEx Super Saver Shipping

Monitors

ViewSonic

$537.00

VP171B 17" LCD Monitor
- SXGA 1280x1024
- Ultra-Slim 0.67" Bezel Design
- Dual 15-pin D-sub, Digital DVI-D Inputs

BenQ

$569.00

FP951
19" TFT Active Matrix LCD
SXGA 1280x1024
Dual D-sub, DVI-D Inputs
Ultra-Slim Bezel Design

SAMSUNG

$447.00

710N-Black
17" a-si TFT/TN LCD
1280x1024 Narrow Bezel
Built-in Power Supply
MagicTune™ & MagicBright™

Intel

intel

New Item

$629.00

Pentium 4 Processor 560
- 3.60 GHz
- 1MB L2 Cache
- Socket LGA 775
- 800MHz FSB
- 90nm Process
- Hyper Threading Technology
- Model# BX80547PG3600E

intel

$221.00

Pentium 4 Processor 530 3 GHz 800MHz FSB
1MB Cache Socket 775 HT
Model# BX80547PG3000E

intel

$181.00

Pentium 4 Processor 520 2.80 GHz 800MHz FSB
1MB Cache Socket 775 HT
Model# BX80547PG2800E

Notebooks

IBM

$1629.00

ThinkPad T41
- Intel Pentium M 1.4 GHz
- 14.1" TFT (1024x768 @ 24-bit)
- 256MB RAM
- 40GB Hard Drive
- 10/100/1000 LAN, Wireless 802.11b
- Windows XP Pro

TOSHIBA

$1429.00

Satellite A40-S270
Pentium M 2.8 GHz
15" TFT (1024x768)
512MB RAM 60GB HD
Wireless (802.11g)
Windows XP Pro

FUJITSU

$1599.00

LifeBook S7010
Pentium M 1.5 GHz
14.1" TFT (1024x768)
256MB RAM 40GB HD
Wireless (802.11a+b/g)
Windows XP Home

Networking

NETGEAR

$44.99

MR814NA
- 11Mbps Cable/DSL Wireless Router
- Smart Wizard Auto-Installation
- 802.11b Wireless Freedom
- Firewall, 128-bit Encryption & Parental Controls

D-Link

$59.99

DWL-G650
High-Speed 2.4GHz AirPlus
XtremeG 108Mbps Wireless
Cardbus Ethernet Adapter

LINKSYS
A Division of Cisco Systems, Inc.

$71.99

WMP54G
54Mbps Wireless-G 802.11g
PCI Ethernet Adapter

Cases

Antec

$135.00

P-160WF
- Silver Aluminum Performance Series
- ATX Mid-Tower Case w/ Side Window Panel
- Swiveling Front Control Panel Features LED Display

TURBO CASE

$69.00

ATXB6KLW-BK/350
X-Infinity Black ATX
Mid-Tower Case w/
350 Watt Power Supply

RAIDMAX

$68.00

ATX-822WBP
Zinc Plated Steel ATX
10-Bay Case w/ 420 Watt
Power Supply

Speakers

Klipsch

$134.99

Promedia GMX A-2.1
- 2.1 Personal Audio System
- 78 Total Watts SWS link for Add-on Subwoofer
- Wired Control MicroTractrix Horn

CREATIVE

$415.00

GigaWorks S750
700 Watts THX Certified
7.1 Speaker System
Two-Way Satellite Speakers

Logitech

$29.00

X-120 2.0 Speakers
Custom Audio Balance
Headphone Jack & Internal
Power Supply

U.S.Robotics

$39.99

USR8003
Broadband Internet Router
5 10/100 Mpbs Ethernet Ports
Universal Plug & Play

Zonet

$29.99

ZEW2000C
11Mbps Wireless USB 802.11b
Adapter Up to 128-bit WEP
2.4 GHz Frequency Band

ALLIED

$69.00

AL-B500E
500 Watt Power Supply ATX
12V w/ Two Aluminum Fans
Ideal for Intel & AMD Systems

Antec

$110.00

LS-968(ARIA)
Aluminum MicroATX Cube
Case w/ 300 Watt Active PFC
Power Supply

FUJITSU

$240.00

MAP3735NP
73.5GB 10,025 RPM SCSI
Ultra 320 68-pin 8MB Buffer
Hard Drive - OEM

SanDisk

$95.00

SDCZ2-512-A10
Cruzer Titanium USB Mini
Flash Drive 512MB Capacity
Hi-Speed USB 2.0 Certified

©2004 Newegg.com Inc. All rights reserved. Newegg.com is not responsible for typographical errors. All typographical errors are subject to correction. All Items are sold as retail unless otherwise noted. OEM products may be sold without packaging, documentation or other articles. All prices and specifications are subject to change without notice. Prices do and will fluctuate dramatically due to the nature of the market. The listed prices were the current price at the time this ad was submitted to the publisher on 6/7/04.
*On These Advertised Specials.

SHOPPERS' COMPUTER SHOPPER 2003 CHOICE

CIRCLE OF EXCELLENCE BizRate

Fulfilling over **10,000** orders a day...and winning awards doing it!

Compiled by
Steve Smith

Ultralights . . . How Low Can They Go?

Ultralight notebooks sank lower in two ways this month with the introduction of a sub-$1,000 unit and a pencil-thin 1.85-pounder. Averatec's 3200, with the AMD Athlon XP-M 2000+, 40GB hard drive, 12.1-inch display, and combo DVD/CD-RW drive, is a mere 4.2 pounds but starts at $999. You even get built-in 802.11g wireless. And while Sony has never been known for bargain pricing, its $2,999 VAIO X505 is slimmer than Lara Flynn Boyle after aerobics. This featherweight squeezes a 1.1GHz Pentium-M and 10.4-inch display into a closed case only 0.83 inches high at its thickest end and 0.38 inches at its thinnest. After years of lugging 7- and 8-pound behemoths through airports, consumers have made ultralights a hot category, so expect thinner and even cheaper models in the months ahead. ▲

The Averatec 3200 series of budget ultralights dips beneath the $1,000 barrier, and the 1.8-pound Sony VAIO X505 keeps 1.1GHz of power in a pencil-thin package.

Servers Get Closer To Home

Face it. With the boom in home networking and gigabytes of digital media just begging to be shared, it is just a matter of time before you're going to need a home server to manage it all. Network kingpin Linksys and storage veteran Maxtor hope to be ready by partnering to develop networked storage solutions for consumers. As a first step, the companies are co-marketing Maxtor's OneTouch External Hard Drive with the new Linksys Network Storage Link. Via its USB 2.0 port, the Maxtor drive connects to the Linksys unit, which then connects by Ethernet to the home router, which can send it wired or wirelessly to the home network. And, yes, we're all hoping that the Linksys/Maxtor marriage results in somewhat more integrated solutions soon. ▲

The new Linksys Network Storage Link can put your external USB 2.0 hard drive on the network.

Epson Prints A Next-Gen Monitor

Epson has used its inkjet printing technology to print out a next-gen flat-screen display. Its 40-inch full-color OLED screen is the largest prototype yet of a technology many experts believe will replace the LCD. Unlike LCDs, which use reflected light, OLEDs are made of self-luminescent organic material layered on thin film transistors. They will use less power and produce brighter, higher-contrast displays that can be viewed in a wider range of lighting conditions than LCDs. Fabricating this sandwich of electronic and organic materials on a large scale has been one of the challenges for moving OLEDs into production; Epson is hoping its inkjet approach will win over competing techniques. (See page 63 of this issue for more on OLED technology.) ▲

Epson's 40-inch OLED display demonstrates how the screens of the future will be bright, lush, and even lighter and more flexible than current LCDs.

A PVR For Your Life

Did you see that? Never ask that question again with the Deja View (www.mydejaview.com) Camwear Model 100. This "wearable camera" clips to glasses or a hat to store the last 30 seconds of its video and audio feeds in a memory buffer. To preserve that slice of life, just press the record button and file the clip in 30fps MPEG-4 format to a memory card in the attached Hip Pack. A 64MB SD card stores over 16 30-second clips. You can transfer the video directly to a PC for burning to disc. This would be a great way to set the record straight on who said what in heated discussions with loved ones . . . if you can get them past the fact that your spectacles have a cam pointed at them. ▲

Hardware Mole

Wi-Fi Bullets

It's official: Everything is about to go wireless, including ammunition. University of Florida researchers have created a "smart bullet" that sends information back to a notebook or handheld device from 231 feet away. This bullet, which is 0.68 inches in diameter, is not designed to harm people but to penetrate potentially hazardous material from a safe distance and report back on traces of explosives. Adhesive on the front of the bullet helps it stick to a target, and an internal battery and transmitter let it communicate with a PC. Unlike most bullets, this one is reusable. ▲

Better Batteries Through Zinc?

Zinc oxide may fuel the battery of the future, according to Intel, which announced a new partnership with start-up Zinc Matrix Power to develop the technology as a replacement for the aging lithium-ion cells. According to ZMP, its prototype zinc oxide batteries demonstrate twice the energy efficiency as current solutions, and they are safer to use. Intel intends to build support for zinc oxide power and recharging schemes into the controller logic of Centrino chipsets. Battery production should start in 2006. ▲

Softball-Sized Computing

Japanese company Personal Media has launched a fully functional PC that is the size of a softball but called Teacube. Using an operating system for embedded systems, the T-Engine Core, and a NEC VR5701 CPU, the 2-inch cube has a full complement of I/O ports and 64MB of RAM. The Teacube connects directly to a keyboard, mouse, and 1,024 x 768 display, and a Web browser and text and script editor are built-in. While we don't think this cube is going to fit in anyone's pocket, it actually weighs only 6 ounces, less than some PDAs. ▲

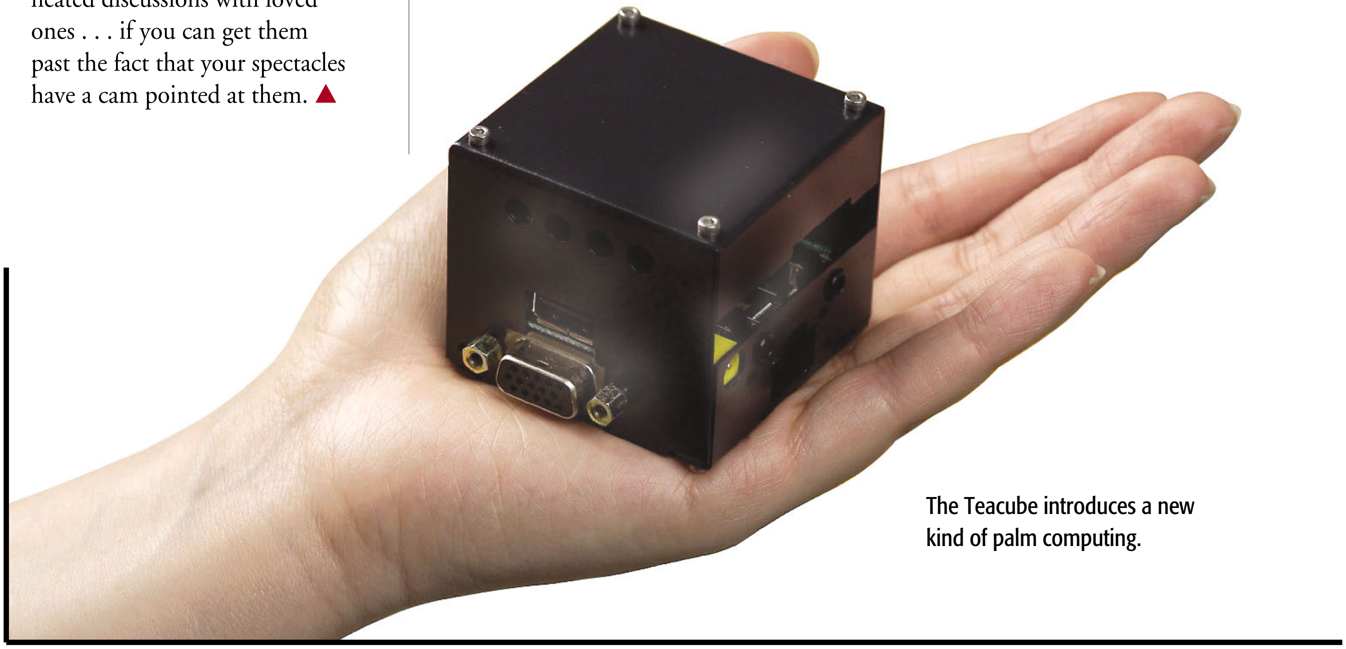

The Teacube introduces a new kind of palm computing.

Compiled by Dean Takahashi

NRAM Memory-Chip Prototypes In The Works

LSI Logic has begun making prototype chips for the 3-year-old nanotechnology startup Nantero. The 1Mb memory chips that the company is calling NRAM (nanotube-based/nonvolatile RAM) use carbon nanotubes—long strands of 1-nanometer-thick carbon that resemble chicken wire—to connect metal electrodes on a mostly standard chip. When opposite charges are applied to the carbon nanotube, atomic-scale van der Waals forces push the wire mesh downward, making contact with metal electrodes. That allows the storage of a digital bit, or 1. When like charges are applied, no contact is made. Greg Schmergel, CEO of Nantero, says that the chips will use several times less power than today's chips, and the cost will be cheap because the company uses standard chip-making processes with just a couple of extra steps. The benefits: The memory is nonvolatile, like flash memory or hard drives, and speedy like dynamic RAM. Schmergel sees no density limit to the memory chips; he believes that NRAM could replace most forms of chip storage today. That could lead to instant-on computers, which would begin operating the moment you touch them. ▲

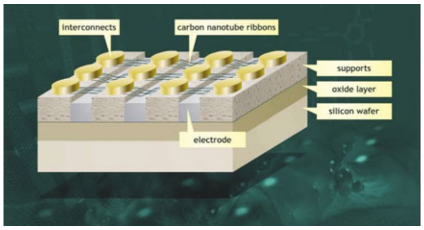

Zarlink Ships A Camera Chip You Can Swallow

Given Imaging is making a camera pill that helps diagnose gastrointestinal disorders. Thanks to an ultra-low-power chip from Zarlink Semiconductor, the pill can capture images as it travels through the digestive tract and beam them back to a doctor. Patients swallow the pill, which has a microchip camera, two silver-oxide batteries, light-emitting diodes, and a radio-frequency transmitter. The pill's RF antenna beams the images back to a data recorder that the patient wears on a belt. The LEDs act as a flash for the camera. The data recorder transfers the images to a workstation, which reconstructs the images of the digestive tract with image-processing software. The images can help doctors diagnose previously undetectable disorders, and the whole process is noninvasive for the patient. Now that's taking digital photography to new depths. ▲

Power Problems For Moore's Law?

Once upon a time, chip engineers could ignore power. But with the latest Pentium 4 chips burning more than 100 watts, the writing on the wall is clear: Pretty soon, the heat from leaking current will be big enough to melt a chip. Leakage, which is the power dissipated even when a chip is idle, increases several times over with each new generation of manufacturing processes that reduce the size of chips. John Kelly, an IBM executive, said in a recent speech that IBM can make an 8-nanometer transistor today, which is about a tenth of the size of current chip transistors. But, he added, "It leaks power like a sieve." Takayasu Sakurai, a professor at the University of Tokyo, says that chipmakers must reduce leakage a hundredfold in the next decade to stay on the path of Moore's Law, which posits that the number of transistors on a chip can be doubled every 18 months. ▲

Watching The Chips Fall

Here is pricing information for AMD and Intel CPUs.

Released	Original Price	Current Price (06/16/04)	Last Month's Price
AMD Athlon XP 2800+ 333MHz FSB			
11/14/2002	$397**	$101	$100
AMD Athlon XP 3000+ 333MHz FSB			
2/10/2003	$634*	$131	$130
AMD Athlon XP 3000+ 400MHz FSB			
5/26/2003	$280*	$147	$150
AMD Athlon XP 3200+ 400MHz FSB			
5/26/2003	$464**	$177	$179
AMD Athlon 64 2800+			
1/6/2004	$193**	$172	$170
AMD Athlon 64 3000+			
9/23/2003	$218**	$211	$211
AMD Athlon 64 3200+			
9/23/2003	$417**	$270	$270
AMD Athlon 64 3400+			
1/6/2004	$417**	$404	$404
AMD Athlon 64 3500+			
6/1/2004	$500**	$500**	N/A
AMD Athlon 64 3700+			
6/1/2004	$710**	$710**	N/A
AMD Athlon 64 3800+			
6/1/2004	$720**	$720**	N/A
AMD Athlon 64 FX-53			
3/18/2004	$733**	$768	$745
Intel Pentium 4 2.6GHz 800MHz FSB			
5/21/2003	$228*	$170*	$174*
Intel Pentium 4 2.66GHz 533MHz FSB			
8/26/2002	$401**	$179*	$173*
Intel Pentium 4 2.8GHz 533MHz FSB			
8/26/2002	$508**	$165	$141
Intel Pentium 4 2.8GHz 800MHz FSB			
5/21/2003	$278*	$167	$182
Intel Pentium 4 2.8GHz 800MHz FSB 1MB cache			
2/2/2004	$178**	$167	$173
Intel Pentium 4 3.06GHz 533MHz FSB			
11/14/2002	$658*	$208	$218
Intel Pentium 4 3GHz 800MHz FSB			
4/21/2003	$417**	$229	$212
Intel Pentium 4 3GHz 800MHz FSB 1MB cache			
2/2/2004	$218**	$215*	$218*
Intel Pentium 4 3.2GHz 800MHz FSB			
6/23/2003	$637**	$275*	$276*
Intel Pentium 4 3.2GHz 800MHz FSB 1MB cache			
2/2/2004	$278**	$278	$278
Intel Pentium 4 3.4GHz 800MHz FSB			
2/2/2004	$417**	$419*	$419*
Intel Pentium 4 Extreme Edition 3.2GHz 2MB cache 800MHz FSB			
11/3/2003	$925**	$843	$845
Intel Pentium 4 Extreme Edition 3.4GHz 2MB cache 800MHz FSB			
2/2/2004	$999**	$989	$1,079*

* Retail price
** Manufacturer's price per 1,000 units
Other current prices, if indicated, are lowest OEM prices available through Pricegrabber.com.

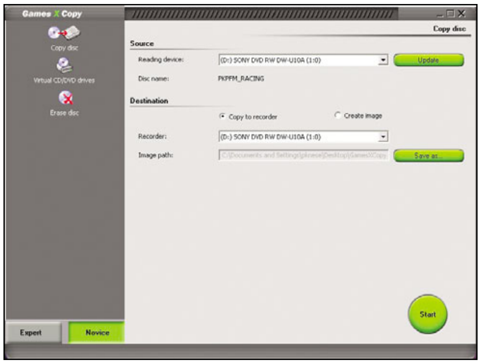

Copy disc

Source

Reading device: (D:) SONY DVD RW DW-U10A (1:0)

Disc name: PKPFM_RACING

Destination

● Copy to recorder ○ Create image

Recorder: (D:) SONY DVD RW DW-U10A (1:0)

Image path:

The game industry's lawsuits against 321 Studios for its game disc copying software may be the last nail in the company's coffin.

321 . . . Stop Copying

The controversial and oft-sued makers of disc copying software 321 Studios made a new enemy recently, the gaming industry. Hoping to end the sale of the company's GamesXCopy tool for making backups of copy-protected PC games, Electronic Arts, Atari, and Vivendi Universal Games filed suit against 321 for violating the 1998 Digital Millennium Copyright Act, which prohibits software that deliberately thwarts copy protection schemes. This suit comes after the movie industry succeeded in getting an injunction earlier this year against further sale of 321's DVD copying software. The organized attacks are having an effect, as 321 admits it has had to lay off hundreds of workers because of decreased sales and has shrunk from 400 to 50 employees. ▲

Half-Life 2 Thieves Get Fragged

One of the most famous software heists in history, the theft of vital code for the hotly anticipated Half-Life 2 game, is closer to being solved. The FBI has made arrests in the case, which apparently involved people from several countries. Half-Life 2 developer Valve Software announced that immediately after the code was stolen in October 2003, the online gaming community rallied to help find the culprits, who also distributed pieces of the code on the Internet. The final game should be out this fall. The FBI says its investigation is not over. ▲

Software Shorts

A Quick Opera . . . By Phone

As all mobile Web browsers know, surfing the Internet on a phone is a special kind of slow torture. Opera's new Mobile Accelerator claims to speed up cell phone Web browsing up to 250% and reduce mobile bills. The tool works on Sony Ericcson P800/900 models and all Series 60 phones by running Web traffic through Opera's proxy server to strip out unnecessary codes and graphics. The service is available for $31 a year, but anyone with a compatible Web-enabled phone can get a 14-day free trial by pointing their cell browser to www.opera.com/proxy. ▲

Apple Plugs A Hole

Windows isn't the only OS with security woes. In what some security experts were calling an "extremely critical" hole in systems running Mac OS X, hackers could execute malicious operations on a Mac. The flaw allowed the use of "URL schemes," where a part of a Web address actually can command your machine to delete its own home directory or issue other arbitrary commands. An initial patch from Apple in May failed to plug this gaping hole, but a subsequent second patch in June apparently does the trick. ▲

What Do You Mean My Phone Is Infected?

File this under "It Was Just a Matter of Time." Security firm Kaspersky Labs says it has detected the first cell phone virus. The Cabir is a worm that can affect phones running the popular Symbian OS and is disguised as a security manager. Written by the 29a crew of virus writers, Cabir does not seem to be malicious but is more of a proof-of-concept worm. It distributes itself via Bluetooth by automatically scanning for nearby phones and sending itself to the first one it hits. ▲

Truly Twisted

For most of us, Twister was the game of absolute last resort reserved only for the rainiest of days—and during blackouts. Now, thanks to the University of California, Berkeley (www.teleactor.net/tele-twister), we can enjoy the classic game without taking our shoes off. Tele-Twister is actually an experiment in "collaborative telerobotics: " how groups of people cooperate to control a machine. In prescheduled live events, you join the blue or red teams, which vote on where human team members at the site place their hands or feet. A Java applet shows a video feed of your team's player on the Twister board and a visual representation of the familiar Twister spinner to show how your team is voting. Your team wins or loses in the usual way, when the opponent falls over, but the online player also gets scored on how quickly and well she anticipated the group's consensus for the best next move during the game. While we're still not convinced Twister is any fun, at least Tele-Twister lets us keep our sneakers on and our faces untwisted from another player's unruly hair. ▲

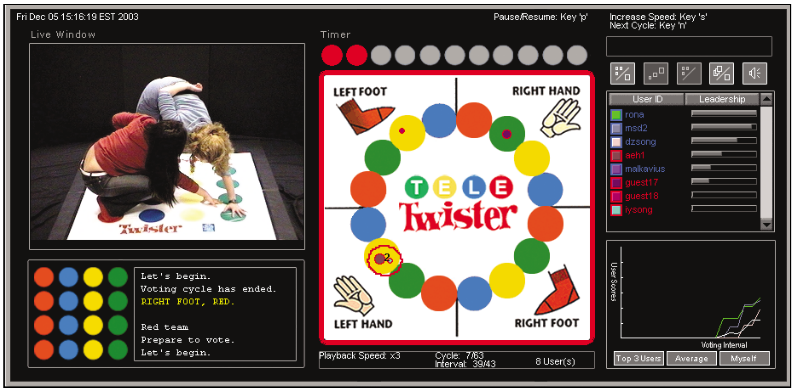

Fri Dec 05 15:16:19 EST 2003

```
Let's begin.
Voting cycle has ended.
RIGHT FOOT, RED.

Red team
Prepare to vote.
Let's begin.
```

Playback Speed: x3
Cycle: 7/63
Interval: 39/43
8 User(s)

Only at Berkeley can Twister become an experiment in "collaborative telerobotics."

SPAM Fans Unite!

As if to prove that there is room for all strange tastes on the Web, Hormel Foods has opened the SPAM Fan Club online (fanclub.spam.com) for lovers of the infamous processed meat product. Yes, sodium nitrate lovers, this is where you will find (we couldn't make this up) SPAM haiku, recipes for SPAM sushi (aka "SPAM Musubi"), and SPAM trivia quizzes. Naturally there is a SPAM newsletter, "A Slice of Spam," and the requisite screensaver, flying SPAM cans, for your Windows or Mac displays. You can even ask the resident guru, SPAMMY, whether there is a national SPAM day. "No, sadly, there isn't," SPAMMY replies. "But there should be." We're going to stop now because it's getting too weird even for us. ▲

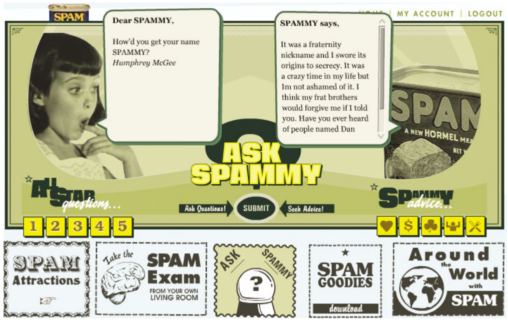

Dear SPAMMY,

How'd you get your name SPAMMY?
Humphrey McGee

SPAMMY says,

It was a fraternity nickname and I swore its origins to secrecy. It was a crazy time in my life but Im not ashamed of it. I think my frat brothers would forgive me if I told you. Have you ever heard of people named Dan

Got a question? At the SPAM Fan Club, a can of SPAM has the answer for you.

The Email Lie Detector

If you want a truthful answer from someone, stay off the phone and try asking him via email. Technology affects lying, says Cornell Assistant Professor of Communications Jeffrey Hancock. In a recent study involving 30 students, he says, "People lied most on the telephone and least in email, and lying rates in face-to-face and instant message interactions were about equal." Apparently having a paper trail in email is one of the motivating factors in telling the truth, as well as the fact that most lying is unplanned and tends to occur as a spontaneous response in a live conversation. In 2,000 conversations among students tracked across a week, 26% of them involved a lie of some kind. Thirty-seven percent of phone conversations involved deception, and 27% of face-to-face exchanges contained a lie. Only 21% of instant messages and 14% of emails contained lies, however. We're guessing that this principle does not apply to those spammers and all of the body parts they claim they can enlarge for us. ▲

Praise The Lord & Click Your Mouse

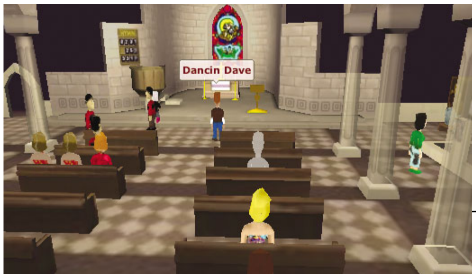

Church of Fools is one chapel where talking in the pews is encouraged.

Can't make church this week? Don't feel guilty; just browse over to Church of Fools (www.churchoffools.com), a 3D online church set up by the Ship of Fools, a Web site devoted to critical discussion of religion. You choose your avatar and enter a 3D chapel where fellow parishioners can mingle or attend prescheduled sermons hosted by noted preachers. There is even a downstairs crypt. About 8,000 people visit every day for this three-month experiment in cyber-worship. The site is intended for those "on the edges of faith," so attendees are warned they can meet up with frank discussion and offbeat characters. The editor, Simon Jenkins, says that he has already encountered virtual Satans, avatars claiming to be the Prince of Darkness. And, yes, there is a donation plate of sorts—an opportunity to contribute via credit card. ▲

New On The 'Net

Did You Get My Email?

A new online service, DidTheyReadIt (www.didtheyreadit.com), lets you track not only whether email you send to others has been opened, but also how long it stayed open and the location from which the recipient accessed the email. Once you subscribe to the service ($49.99 a year), you simply append ".didtheyreadit.com" to the emails you send, and you will receive a message indicating when, where, and for how long it was opened. Unlike the Read Receipts function in Outlook, this approach is invisible to the recipient. ▲

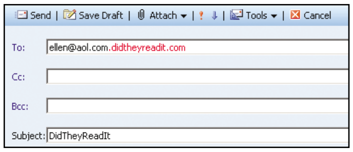

To: ellen@aol.com.didtheyreadit.com
Cc:
Bcc:
Subject: DidTheyReadIt

By adding a small command to an email recipient's address, you can have DidTheyReadIt track how long someone read your message.

Seeing Starz In Broadband

The downloadable media model finally comes to the film world with Starz Ticket (www.starz.com), the first broadband movie subscription service. This partnership between RealNetworks and Starz, the premium cable company, lets subscribers stream or download as much as they like for $12.95 a month from a rotating library of more than 100 recent and classic films. Initial titles include "Finding Nemo" and "Hulk." While these downloadable movies won't play on your iPod just yet, you can run them anytime on your PC or via a TV-Out card to your TV. ▲

Candidates On Demand

For those who just can't get enough election year stump speeches (and we know there are some of you out there), CampaignSearch (www.campaignsearch.com) is your own personal Google. You can search video of the candidates from Kerry and Bush sites, as well as C-Span, PBS, and a host of other sources. The clips are indexed and cut into manageable clips around the keywords you plug into the search box. ▲

Compiled by Kevin Savetz

BIOS Upgrades Available Online

Before you send another motherboard to the landfill, consider upgrading the BIOS and giving your PC a new outlook on life. Here are a few recently released upgrades. Readers can check out www.cpumag.com/cpuaug04/bios to see our entire upgrade list.

Motherboard	Date Available	URL
ABIT KV8-PRO	06/02/2004	www.abit-usa.com/downloads/bios/bios_revision.php?categories=1&model=176
ABIT VI7	06/09/2004	www.abit-usa.com/downloads/bios/bios_revision.php?categories=1&model=133
AOpen AX4C Max II	06/10/2004	download.aopen.com.tw/downloads/userdownload_List.asp?Publish=Yes&RecNo=8114&Model=AX4C%20Max%20II
AOpen MK89-L	06/08/2004	download.aopen.com.tw/downloads/userdownload_List.asp?Publish=Yes&RecNo=7928&Model=MK89-L
ASUS P4S800D	06/11/2004	www.asus.com/support/download/item.aspx?ModelName=P4S800D
ASUS P4S800D-E Deluxe	06/11/2004	www.asus.com/support/download/item.aspx?ModelName=P4S800D-E%20Deluxe

Compiled by Steve Smith

If you have an electrical engineering degree and know PC architecture backward and forward, you could help launch the next generation of gaming consoles. The Phantom game deck, which is expected to ship in November, will stream PC games to your TV via a broadband connection, and maker Infinium Labs needs a Senior Electrical Design Engineer to help design the innards. You will develop and maintain product specs as the unit goes to production, evaluate components and suppliers, and define the electrical design process along with an offshore development team. Fly or flop, the Phantom project is sure to be one of the most discussed high-tech items this Christmas. Either you will be in on the ground floor of something big or you'll be able to brag someday about having been part of one of the great busts of gaming's go-go era. ▲

RAW DATA

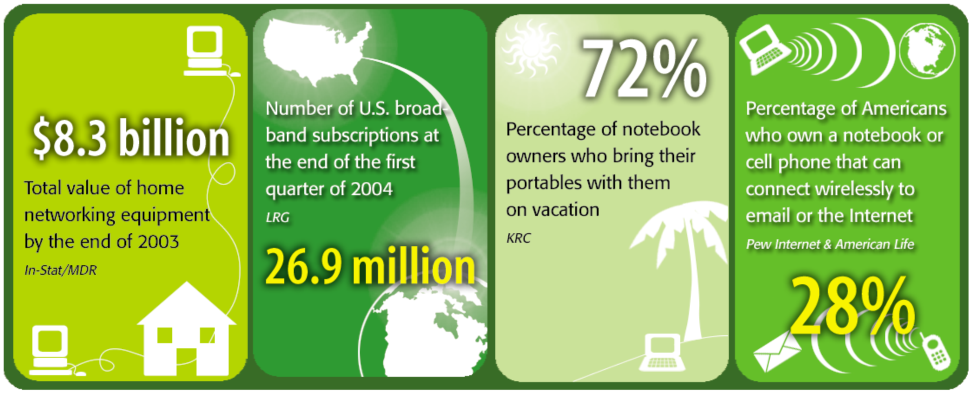

$8.3 billion
Total value of home networking equipment by the end of 2003
In-Stat/MDR

Number of U.S. broadband subscriptions at the end of the first quarter of 2004
LRG
26.9 million

72%
Percentage of notebook owners who bring their portables with them on vacation
KRC

Percentage of Americans who own a notebook or cell phone that can connect wirelessly to email or the Internet
Pew Internet & American Life
28%

PVRs Come Into Their Own

Considering the high profile of the TiVo and other PVRs, you would think that half of us already own one. In fact, until 2003, the PVR was among the most overhyped and undersold digital devices around, with only 1.5 million units shipped worldwide in 2002. But that is about to change, says InStat/MDR analyst Mike Paxton, as combo DVD/PVR decks became very popular in Japan last year, and U.S. pay-TV providers start incorporating PVR functionality into set-top digital cable boxes. PVR sales tripled in 2003 (4.6 million units) and will more than double again in 2004 (11 million) with 6.5 million in the United States.

Mobile Media Usage (U.S. vs. UK)

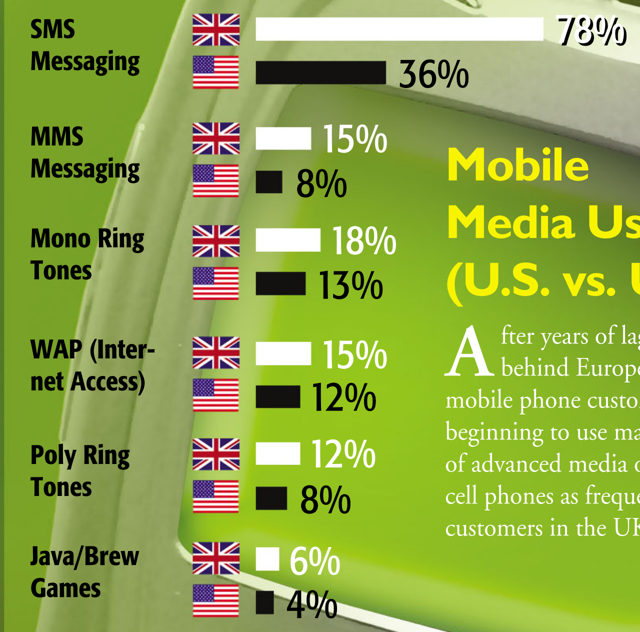

After years of lagging behind Europe, U.S. mobile phone customers are beginning to use many types of advanced media on their cell phones as frequently as customers in the UK.

Source: Enpocket

OCZ Technology

www.ocztechnology.com

"GOT EB DDR?"

ENHANCED BANDWIDTH

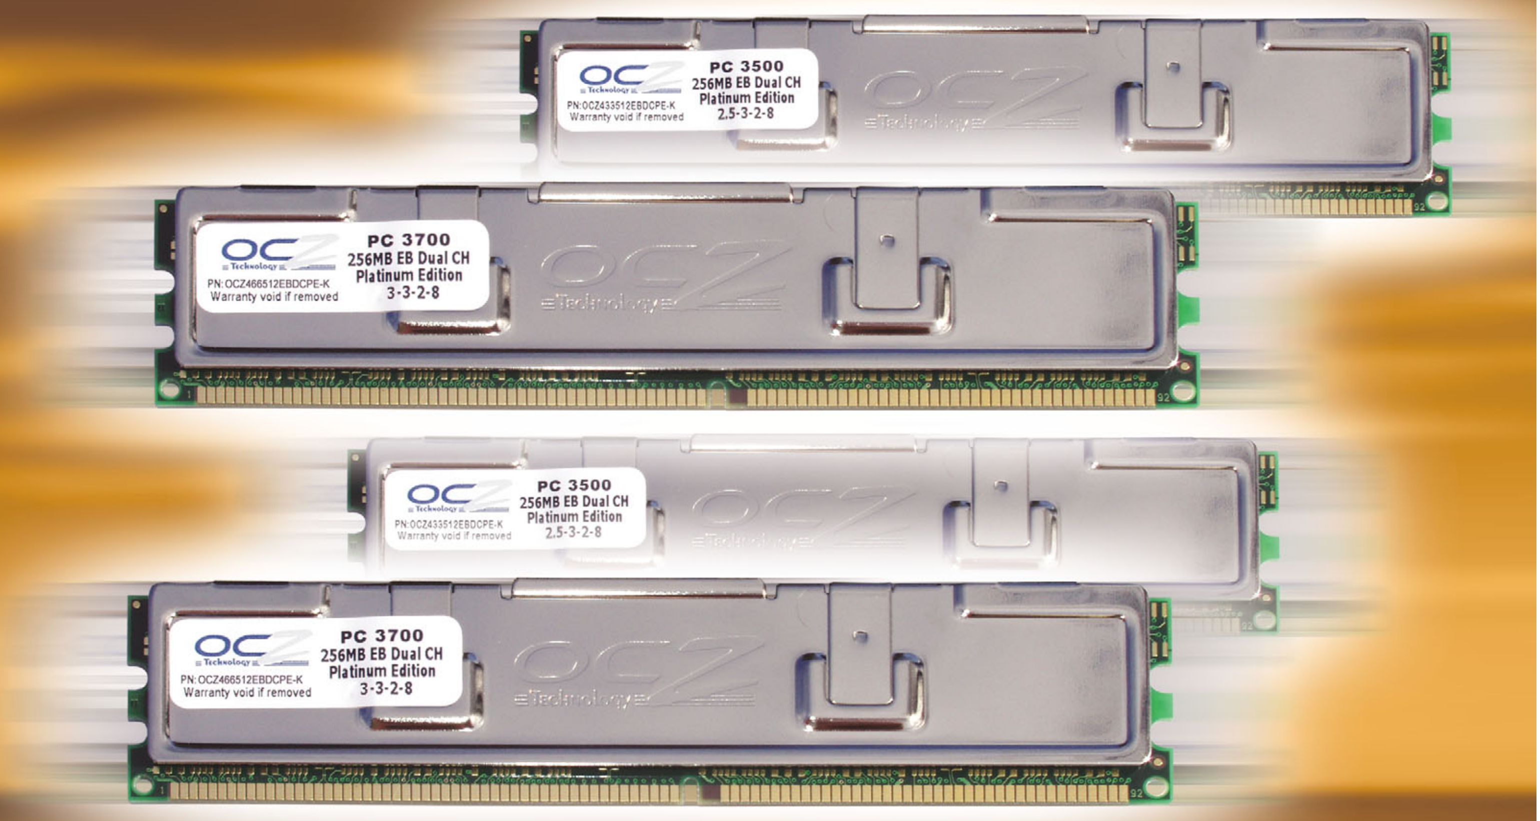

SPECIFICATIONS

PC3200 DDR400 CL 2.5-3-2-8
PC3500 DDR433 CL 2.5-3-2-8
PC3700 DDR466 CL 3-3-2-8

512MB (2x256MB) and 1GB (2x512MB) Dual Channel Kits

Mirrored Platinum Copper Heatspreader
2.8Volts
Lifetime Warranty

DESCRIPTION

Enhanced Bandwidth (EB) technology is a means of increasing memory bandwidth through the optimization of memory latencies for the best possible interaction between the system memory and the chipset and memory controller. It uses low tRP and tRCD latencies in combination with a Variable Early Read Command to allow for the highest possible effective data bandwidth. In most applications, the 2.5-3-2 (CL-tRP-tRCD) will deliver bandwidth that is indistinguishable from CL-2 modules.

OCZ Technology
295 Santa Ana Court, Sunnyvale, CA 94085 USA
http://www.ocztechnology.com
(408) 733-8400 Phone (408) 733-5200 Fax
Sales Team: (408) 733-8400 ca@ocztechnology.com

newegg.com

newegg.com

Axion Tech

axiontech.com

ATACOM.com

atacom.com

eXcaliberPC.com

excaliberpc.com

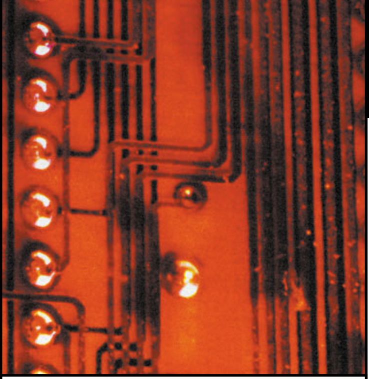

by Alex St. John

What's In A Qubit?

Alex St. John was one of the founding creators of Microsoft's DirectX technology. He is the subject of the book "Renegades Of The Empire" about the creation of DirectX and Chromeffects, an early effort by Microsoft to create a multimedia browser. Today Alex is President and CEO of WildTangent Inc., a technology company devoted to delivering CD-ROM quality entertainment content over the Web.

While we're overclocking our CPUs to get an extra 5% boost, there are folks out there working on the kind of computer technology it will take to run a "Star Trek" Holodeck one day. I'm talking about quantum computing.

Everything you may have learned in high school and college chemistry was probably wrong, which is why quantum computing doesn't seem completely obvious and comprehensible to everyone the first time they hear about it. Albert Einstein himself went to his grave incredulous of the ramifications of quantum mathematics. Nonetheless, nearly all modern computing technology is based on veracity of the exact same mathematical foundations.

Typical chemistry classes teach us that the atomic world is made out of tangible little objects that behave like the objects we experience every day. For example, when I lose my car keys every morning, they're usually somewhere in my house. If I look on my dresser and find them, then the keys were on my dresser and there was zero actual probability that I would have found the keys anywhere else I looked.

. . . a qubit stores probabilities of being found in a particular state, not the state itself.

When a quantum physicist loses his car keys, he thinks of the keys as actually being in many probable locations all over the house at the same time. If he looks on the couch and finds the keys he doesn't conclude that the keys must have been there all along. Instead, he concludes that the "wave function" describing the location of the keys collapsed when he looked at the couch forcing the keys to exist there when he looked, thus eliminating them from the dresser where they might otherwise have been found, as well. Until he looks, the keys are everywhere at once. Sounds nutty, but the problem is that for a quantum physicist this is the way that the world actually behaves, at least on an extremely small scale. Adopting this strange way of thinking is key to understanding quantum computing.

A classical computer stores information in units called bits. A bit can be found in one of two states, on or off, usually represented by one and zero. Stringing bits together allows us to store more complex units of information like this column. Our quantum physicist has a "bit," as well, but it's called a qubit. A qubit can store an infinite amount of information because a qubit stores probabilities of being found in a particular state, not the state itself. A bit can store whether or not the car keys are on the dresser, yes or no. A qubit can represent all the probabilities of where the car keys might have been in the house before our physicist looked for them. This is an exotic but powerful idea.

Qubits that store infinite information but only produce an on or off result when examined are at the heart of quantum computing, but what use is infinite memory that erases itself any time we read something out of it? The information isn't very valuable unless you can compute against that infinite data and the information you can see is the result of that computation. This is precisely what it appears a quantum computer can do. When I look for my car keys I have to search one location after another to find them, eliminating one location after another until all that is left is the only place the keys could be. Our quantum physicist can effectively search every location in the house simultaneously. If the keys were lost anywhere on the planet Earth, he would find them just as quickly, whereas I might spend a lifetime searching one location at a time. Thus with infinite data storage potential, also comes infinite parallel processing capability.

Where does this data and processing power come from? The answer to that question is still a mystery, but our best modern theories describe a quantum world a bit like the Holodeck in "Star Trek," which can appear to contain an arbitrarily large world in a small space. If all this talk of quantum computing's got you excited, please turn to page 65 for more of it.

Next month, I'll tell you how Intel's QuPentium processor shipping sometime in 2050 might actually be made. ■

Jump the distance and send feedback to TheSaint@cpumag.com.

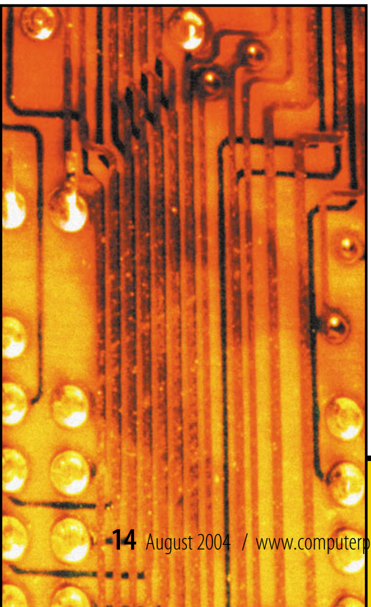

Dream Hardware

Ever get the feeling that you are part of the reason why so many people are against cloning? If all you want is a Mini-You to dress up and boss around, try a Nuvo robot instead. When you're ready for bigger entertainment, even one of those blatantly anti-Orc propaganda films that have been so popular lately, Digital Projection and Samsung are there to put you in the zone.

by Marty Sems

ZMP Nuvo

Oz never did give nuthin' to the Tin Man that he didn't already have. However, robot cobbler ZMP seems to have given its tin simulacrum a bit of a duality complex. Nuvo has the broad-shouldered, hulking mien of a BattleMech, yet he has big, friendly beach balls instead of hands. He can be a stalwart home sentry you can control through your cell phone, yet he's been known to wear rainbow paint schemes that would look great on Mr. Kotter's apartment wall. Nuvo even has a line of clothing and shoes so you can dress him up, for cryin' out loud. This $30,000 fellow has been clumping around now for a while, but we're still awaiting his mass production (and price drop to about $4,600) by the end of the year (www.zmp.co.jp). In the meantime we sincerely hope he sorts out whatever issues he's struggling with.

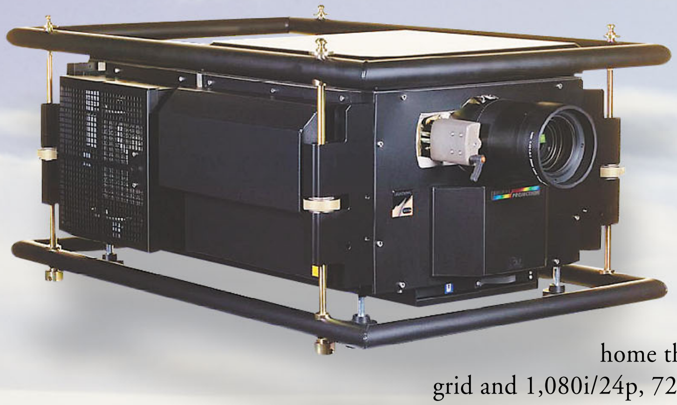

Digital Projection Lightning 35HD

How many times can lightning strike at Digital Projection? As many times as it wants to, we say. The 35HD is the latest and greatest in the Lightning line of uber-projectors ($119,995; www.digitalprojection.com). If you have 4,000 watts to feed it, the Lightning 35HD has 16,000 ANSI lumens to burn into your home theater wall. With a 2,048- x 1,080-pixel grid and 1,080i/24p, 720p, and 480p resolutions on tap, you shouldn't have any trouble finding Nemo in all that screen space. Of course, the Lightning pumps out a scary amount of heat and fan noise, so maybe you should skip the Pixar flicks and save the projector for calamitous disaster movies. The dialogue is pretty immaterial in those, anyway.

Samsung LTP468W

Stepping back from the frontier of huge displays you can't buy at any price, we find the world's largest flat-panel LCD TV currently on the consumer market, according to Samsung (www.samsungusa.com). The 46-inch LTP468W offers HDTV resolution up to 1,080p, as well as 24-bit color depth and a mongoose-quick 12ms pixel response time. Samsung fortified the TV with a decent 800:1 contrast ratio, plus a 30-watt JBL sound system with SRS TruSurround XT audio processing. Look for the LTP468W this month for about $10,000, and keep an eye on the horizon for the company's 57-inch LCD, too. It will be the glimmering thing in the shadow of Samsung's 80-inch plasma prototype.

The Search Continues

Looking For The Perfect HTPC Case

Minuet
$79
Antec
www.antec-inc.com
CPU CPU CPU CPU

In *CPU*, we've looked at several variables of the HTPC (home theater PC) equation. Now with Windows XP Media Center Edition due in an OEM version by year's end, the need to get serious about building a living room box grows larger by the week. We decided to tackle the chassis problem and examine several of the best cases for the task.

This is tougher than it seems. Tower systems are simply aesthetic nightmares, and even SFF boxes, such as the XPC and iDEQ, can seem out of place. A proper HTPC chassis should look natural on a home theater shelf and accommodate industry-standard motherboard form factors. The following are five models for turning your living room into a multimedia PC haven.

Antec Minuet

The Minuet is perfectly named. This MicroATX case is small (3.8 inches high x 12.75 inches wide x 16.8 inches deep), elegant, and, like the classical 17th century French minuet dance, it's designed for couples. The Minuet gleams with a lacquered "piano black" finish on the sides and a brushed-aluminum face. Rubber vibration absorbers and a quiet 220W power supply help hush the chassis, and its narrow width makes it look at home among your other home theater components.

Like many MicroATX cases, the Minuet only accepts half-height cards (so kiss your top-end gaming dreams goodbye), and you'll probably end up with a USB-based TV tuner product. However, Antec does not use a door to conceal its two USB, one FireWire, and two audio ports. Yes, the ports disrupt the bezel's clean lines, but not obtrusively so, and if you're going to use the ports, the end result is cleaner-looking without a rickety door hanging open.

Under the hood, the Minuet is full of convenience. Board stand-offs are part of the chassis floor, not nuts you need to screw in. There's only one drive cage (that rotates up and lifts out), into which you can install two external drives (one 5.25-inch and one 3.5-inch), plus one internal 3.5-inch drive. Antec doesn't adopt screwless mountings, which instinctively bugs me, but rationally I know that HTPCs aren't likely to undergo a lot of drive changes. Front-bezel port cables are the perfect length, the interior is tight without being cramped, and the only annoyance I had was popping out the stiff bezel drive plates.

A narrow vent runs along the top panel above the PCI slots. There's also a spot to mount a 40mm case fan on the rear panel and an 80mm case fan mount on the right panel. This gives you several cooling options, although a 3GHz processor running on an IGP chipset with one hard drive probably has no thermal worries.

Some reviewers have complained about the Minuet's lack of a reset button, but is it really so hard to hold the power button for five seconds? I'll take a cleaner look over more buttons. Antec offers a three-year warranty, precisely because the Minuet is built like a tank and has few things that can go wrong. Of these cases, the Minuet has the best universal appeal as an HTPC chassis.

AOpen H360

Bad news first: This beige chassis with blue and gray face panels was built for corporate desktop use, not living rooms, and in a lineup of sexy black and silver components, it clearly stands out as the ugly duckling.

Fortunately, the rest is mostly good news. The H360 is a well-ventilated FlexATX/MicroATX chassis with a 250W PSU. The dimensions are 13.2 inches x 3.9 inches x 15.95 inches, making it admirably compact for the flexibility it offers. The only drawback is that the slim design only allows for half-height cards.

The main drive cage rotates upward and lifts out for easier drive installation. The really cool innovation is a removable bracket piece that mounts inside the cage. Out of the box, the cage holds one 5.25-inch drive (external) and two 3.5-inch bays (one external). However, if you flip the bracket 180 degrees and pop out a bezel plate, the cage holds two 5.25-inch external bays. Below the cage, secured to the chassis' floor, is a smaller 3.5-inch internal drive cage.

AOpen also supplies a toolless drive mounting kit. The bays are constructed such that drives snap into them fairly snugly already. You can then either secure

CPU Ranking: 0 = Absolutely Worthless CPU CPU C 2.5 = Absolutely Average CPU CPU CPU CPU CPU 5 = Absolutely Perfect

them with screws or snap on a plastic bar to one outer edge of the bay. The bar clips into place and has plastic dowels that run through into the drive's screw holes to lock it in place. A similar toolless clip replaces the backplane slot screws.

The front panel features headphone, mic, and two USB extension ports. There's a hole for a FireWire port option, and another $20 gets you AOpen's side-blowing CPU heatsink fan, which pushes hot air straight into the PSU's intake fan. AOpen uses an oversized 90mm fan in its 250W power supply, yielding a whisper-quiet 24dB noise rating. And because there are no case or GPU fans, this is likely to be one of the quietest, best-priced HTPC cases you'll encounter.

H360
$65
AOpen
usa.aopen.com

Chenbro Hornet

Some HTPC cases are efficient; others are sexy. Few can actually be called "cool," but Chenbro's Hornet is exactly that. Styled to primarily be a MicroATX minitower, the Hornet serves just as well on its side, although this makes it harder to see the hornet head-shaped window in the side/top panel. The bezel has two black doors for the 3.5-inch external bay and FireWire, audio, and two USB ports. A long door conceals the case's only 5.25-inch bay. Both long bezel panels serve as stylized "hornet eyes" (when the case is upright) and are grilles for excellent interior airflow.

Measuring 13.6 inches x 5.4 inches x 13.8 inches, nearly all the Hornet's assembly is optionally screwless and quick. The ODD and HDD bays lock into place with spring clips (you can add screws if you want), and the bays swing up and lift out of the case. With the ODD and 3.5-inch external drives, you merely plug a vaguely arc-shaped shim into two screw holes on the drive's left side and shove the drive into its bay where the shim securely snaps it into place. It takes about 10 seconds to install a drive. The primary HDD bay also swings up, but two screws that run through rubber grommets secure the drive to help reduce vibration. Another internal bay under the ODD converts from 3.5-inch to 5.25-inch if you remove the premounted bracket.

Solid and equipped with a 270W Fortron Source PSU, I was ready to call it a day and take the Hornet home permanently when I made a critical discovery: There was maybe a quarter-inch of clearance between the edge of the 5.25-inch drive cage and a chip mounted on our 5950 AGP card. (The Hornet supports full-height cards.) Just a little pressure from a cable bringing the two into contact could short out the system. I actually saw sparks fly a couple of times. A little foam in the right spot fixes the problem, though.

Some users may find the Hornet too visually loud, but if you like some attitude, Chenbro delivers nicely.

Hornet
$90
Chenbro
www.chenbro.com.tw

Cooler Master Cavalier 2 (CAV-T02)

If you have an ATX board that you want to make your HTPC's foundation, the Cavalier 2 is one of the best products to make it happen, but at the expense of 16.5-inch x 17.2-inch x 5.3-inch dimensions. Cooler Master's impeccable construction has almost no give when you twist the chassis, the black aluminum bezel and swing-down door are thick aluminum, and impeccable machining means no stuck cover or screws that won't go in.

What originally caught my eye was the Cavalier 2's Musketeer Audio Sound Pressure device in the front panel; it's essentially an old-fashioned VU meter like those with a home stereo amplifier. The meter uses a quarter-inch pass-through mounted into a PCI backplane slot. The pass-through has one jack in the back that connects to the VU meter's back and two jacks in front. One external jack connects via a patch cable to the system's stereo-out port and the other to your headphones or speakers. Sure enough, the blue-lit meter does a fine job of twitching along with the PC's audio output. Unfortunately, my meter did just as much twitching whenever there was any hard drive or optical drive activity.

I'm not sure why Cooler Master advertises the case as having tool-free assembly because the only part of assembly that did not require a screwdriver was the snap-in, two-bay (external) 3.5-inch drive cage. The chassis also sports two external 5.25-inch bays and another internal 3.5-inch bay.

Don't let these nitpicks dissuade you. I really like the unit's sleek, curved facing. The 300W, temp-controlled power supply sports plenty of leads, including two SATA connectors. The side-mounted 80mm and twin rear exhaust 60mm case fans are virtually silent, and I really dig that Cooler Master put a volume dial on the front panel below the audio, FireWire, and USB jacks. This unit is tough, graceful, and spacious enough to handle all of your HTPC add-on gear.

SilverStone Technology SST-LC04

First, special thanks to Monarch Computer for generously supplying this unit when SilverStone didn't respond to review queries. Unlike the full-height

CPU Ranking: 0 = Absolutely Worthless 2.5 = Absolutely Average 5 = Absolutely Perfect

Cavalier 2 (CAV-T02)
$160
Cooler Master
www.coolermaster.com

LC01 and LC03, which Monarch also sells, the LC04 is a half-height slim desktop made just the right width to match your stereo components. Unlike other slim-case vendors, Silverstone includes one PCI and one AGP riser so that you can mount one or two full-height cards parallel to the motherboard.

The LC04 comes in black or silver and measures 16.5 inches x 4.5 inches x 12.2 inches, yet accommodates ATX, MicroATX, and Mini-ITX motherboards, making it probably the smallest full-ATX chassis I've seen. The trade-off is drive capacity. The case holds one 5.25-inch drive (external) that mounts onto a combination tray-crossbar that gives the unit much of its rigidity. The single 3.5-inch drive "bay" is actually a shim that juts out from the case's side. One side of the drive mounts against the shim and the other to the chassis' floor. Squeezing a drive into this spot was difficult because the shim sits a pinch too low, and I don't envy the case vibration you'll face if you have a rickety drive.

SST-LC04
$169
SilverStone Technology
www.silverstonetek.com

Similarly, the crossbar actually sat on the heatsink bolted to the back of my NVIDIA 5950 test card. Maybe this will help with heat dissipation, but an extra quarter- to half-inch in case height could have been helpful. I had to remove the backplane shim on my graphics card to get it to sit in the slot.

There isn't enough peripheral capacity here to challenge the LC04's fairly quiet 240W power supply. There are two side grilles for cooling, but only one (by the hard drive) has mounting holes for an optional 80mm case fan, which will leave your interior with an air pressure imbalance. Still, because the CPU fan blows up into a top panel vent grille and my 5950 card fan blows out the backplane slot, I had no thermal problems.

Cosmetically, the LC04 is innocuous and attractive. Two long bars dominate the front panel on my silver sample, with the bottom being a mirror in which sits the power button and the top being a brushed-aluminum door that conceals the optical drive, two USB ports, headphone, mic, and FireWire jack. The FireWire cord ends in a standard 6-wire FireWire connector, not the usual motherboard pin connector, leaving me wondering where I was supposed to plug it.

Despite its drawbacks, the LC04 is still a good chassis for a light-use HTPC that won't need much future expansion.

Last Word

Here's the truth: I have yet to find the perfect HTPC case. Perhaps it will ride in on the coming wave of BTX designs. If I could take Antec's exterior, AOpen's bay flexibility, Cooler Master's VU meter, Chenbro's easy assembly, and Silverstone's allowance for full-sized cards in a half-height case and mesh them together, the perfection would be overpowering. That's not to say these cases aren't good enough to house your HTPC, but I expect better things to come as vendors pay more attention to this niche's needs. ▲

by William Van Winkle

The Case For A HTPC Case

	uATX & ATX capability	Vertical display capability	# of backplane slots	Slot height	PSU	5.25-inch bays (exterior)	3.5-inch bays (interior/exterior)	Toolless
Antec Minuet	No	Yes	4 half-height	Half	220W	1	1/1	No
AOpen H360	No	No	4	Half	250W	1 or 2	1/1 or 1/0	Yes
Chenbro Hornet	No	Yes	4	Full	270W	1 or 2 (1 interior; 1 exterior)	1/1 or 0/1	Yes
Cooler Master Cavalier 2	Yes	No	7	Full	300W	2	1/2	No
SilverStone SST-LC04	Yes	No	2	Full	240W	1	1/1	No

CPU Ranking: 0 = Absolutely Worthless 2.5 = Absolutely Average 5 = Absolutely Perfect

The Award-Winning MSI K8N Neo Platinum Motherboard is Now Available for Purchase!

Experience the 64-bit Computing TODAY!

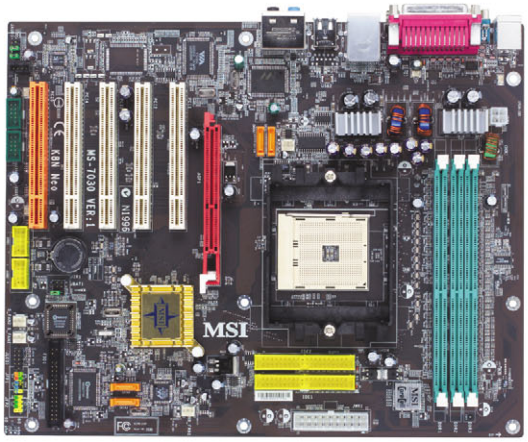

"If you've been holding out for a better Athlon 64 solution, the K8N is definitely the answer you've been waiting for!!!"
- PC Enthusiast Magazine

"The MSI nForce 3 250Gb took every single synthetic and gaming benchmark tested, leaving us with no doubt as to the new performance leader in the Athlon 64 arena ."
- vr-zone.com

"The MSI K8N Neo Platinum is a truly excellent motherboard that competes in every way with the best of the Socket 754 motherboards."
- Anandtech.com

"It's easy to see that the MSI K8N Neo-FIS2R Platinum is a winner . It outperforms the most popular socket-754 boards on the market."
- GamerDepot.com

"The MSI K8N Neo-FIS2R Platinum is an excellent board to work with …… the performance of the MSI K8N Neo-FIS2R Platinum motherboard was excellent. It's easily the fastest Athlon64 board we've tested to date."
- PCStats.com

"Overall, the MSI K8N NEO was a pleasure to work with and extremely fast to use. From what I have seen, I have no problems with recommending the MSI K8N NEO to your short list when shopping for Athlon64 motherboards."
– HardOCP.com

Also Available Athlon 64 Models:

K8TM-ILSR

K8MM-ILSR

K8T Neo-FSR

K8T Neo-FIS2R

MSI
WWW.MSICOMPUTER.COM
Link to the Future

MSI makes the best AMD Platform motherboard.

Contact Us:

MSI Computer Corp.
901 Canada Court
City of Industry, CA 91748
Tel: 626 913 0828
Fax: 626 913 0818

Intel 775 Platform

915P/G & 925X Chipsets Have Arrived

915P/G and 925X
$150 to $250
Intel
(408) 765-9904
www.intel.com

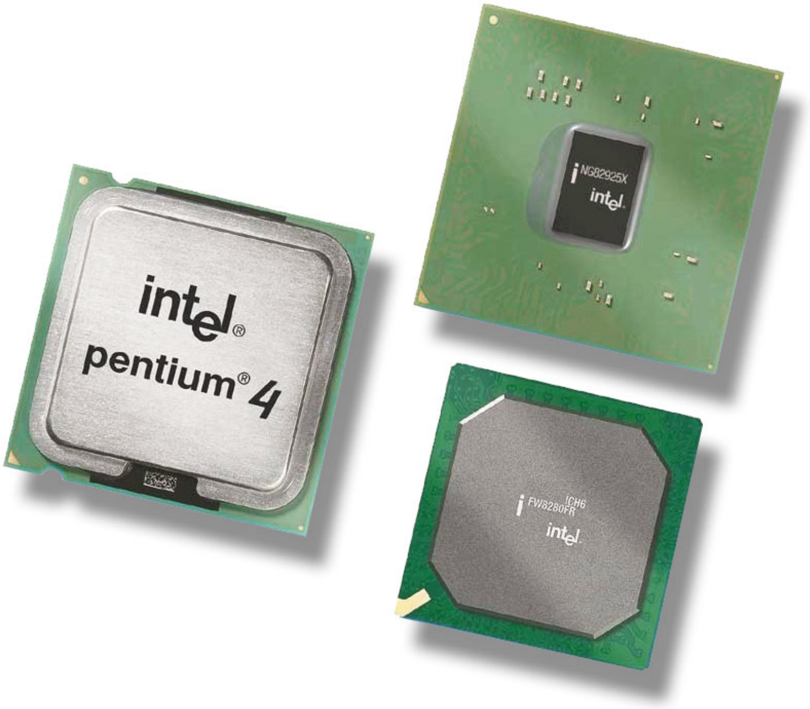

It's been quiet on the Intel chipset front for quite awhile, but with the new 915P/G and 925X chipsets, we are finally starting to see the arrival of PCI Express and DDR2 support. Based upon Intel's Alderwood and Grantsdale chipset, the 775 platform sees the company's biggest shift in years.

It's out with AGP and in with PCI Express, with its full 16-lane graphics connection and 8GBps of bandwidth (actually 4GBps each way when bidirectional). That's a hefty hike when compared to AGP, but games today still don't touch the limits of AGP 8X, so it'll be awhile before we start to tap into the PCI Express ceiling. Even though the power budget has been increased, you'll still be plugging external Molex connectors into those high-end video cards. PCI Express isn't just about graphics, though. Its single-lane I/OC connectors, which provide a point-to-point link, increase the throughput to 2GBps each way (double the old PCI standard).

What's new. DDR memory has been the standard for some time across all platforms, but Intel has decided that now is the time to step up to DDR2. It's still early days, but DDR2 does function at faster clock speeds. There are other advantages, such as a lower operating voltage (1.8V instead of 2.5V), a larger prefetch (four instead of two), and on-die determination for better signal integrity.

Faster clock speeds are one thing, but at this stage the latency increase means that currently memory sticks work with a CAS 4 latency, which isn't the same as the CAS 2 we're used to in the performance world with DDR. This situation will change in the future, but it certainly hampers overall performance of the 9xx series, unfortunately. So we used 533MHz DDR sticks set to 4-4-4-12 timings.

High-definition audio is another new feature (code-named Azalia), which is a considerable improvement over the AC'97 standard. Intel's HD audio is able to do eight-channel, 192KHz, 24-bit audio streams. Multistreaming support allows for the routing of different audio sources to multiple output devices, which is nifty.

Serial ATA isn't new, but for the 9xx series of chipsets, Intel has dubbed its storage capabilities as "Intel Matrix Storage Architecture." The company has upped the number of SATA ports to four and now added RAID 1 functionality. Matrix Raid gives you the performance advantages associated with striping RAID 0 but with the added backup of using RAID 1 with only two drives. Intel's software is extremely easy to use and should help bring RAID even to Ma and Pa's PC.

Onboard graphics aren't something *CPU* readers usually care about, and suffice it to say that the new integrated Intel graphics core won't be winning any awards. It does at least up the ante to a DX9-class graphics chip at the low end. Clocked at 333MHz and with a four-pixel-pipe architecture, it's at least more "extreme" than the last core.

Hardware. On the physical hardware side, the socket 478 is replaced with a new socket, the LGA (land grid array) 775. The extra pins do serve up more power but require a bit of extra care when mounting your CPU. Instead of the (bendable) pins being on the processor itself, they are now on the motherboard socket, with the CPU having flat contact points. You then apply a bit (and we can't stress that enough) of pressure and, with a lever, lock down the CPU.

Specs: PCI Express, DDR2, and HD Audio mainboard for Pentium 4

CPU Ranking: 0 = Absolutely Worthless 2.5 = Absolutely Average 5 = Absolutely Perfect

This new socket means new 775 LGA processors, and along with the new chipset release, Intel is set to release some new CPUs, which also means the debut of the company's new naming scheme. The Prescott P4 is up to 3.6GHz, dubbed the 560, and there is also room for a P4 3.4GHz Extreme Edition. A 3XX series will also show up, denoting the re-release of a budget "Celeron" line of processors.

The performance differences at this early stage are negligible to none, which may dissuade you somewhat from stepping up at this point. However, the number of new features you get to play with is seriously worth your consideration. By the time you're done playing with all of them, maybe, just maybe, game developers will start tapping into their raised PCI Express ceilings and/or memory manufacturers will start nailing down better latencies than we're used to with DDR. Expect the usual OEM suspects from Taiwan to start shipping 915/925 boards over the coming months. (We've already received boards from ASUS.) ▲

by Alex "Sharky" Ross

Processor Benchmarks

	3.4GHz P4EE/925X	3.4GHz P4EE/875P	Athlon 64 FX-53
PCMark 2004			
CPU	5220	5267	4511
Memory	5375	5421	5534
3DMark03	17653	17428	17028
Splinter Cell	76.7	78.2	92.4
Comanche 4	73.5	74.6	78.1
UT2003 Botmatch	72.3	73.6	90.2

AMD Athlon 64 3800+ & FX-53

Athlon 3800+
$720
AMD
(408) 749-3060
www.amd.com
CPU CPU CPU CPU

FX-53
$799
AMD
(408) 749-3060
www.amd.com
CPU CPU CPU C

The Athlon 64 has been winning the hearts and minds of performance enthusiasts since its introduction last year. However, the different sockets and registered memory requirements are still considered to be glaring weaknesses. With the new socket 939 Athlon 64 processors, AMD is taking care of these weaknesses with one big swipe.

The socket 939 platform serves up top-to-bottom-scale Athlon 64 technology on the same platform, with dual-channel memory support across the board. And even better news for FX processor users is that the new platform allows for the use of unbuffered memory, which is cheaper, faster, and more readily available. The top-tier processors that will be available by the time you read this are the Athlon 64 3800+ and the FX-53. The cores remain the same, but the clock speeds are set to 2.4GHz for both, with the latter obviously having double the cache at 1MB.

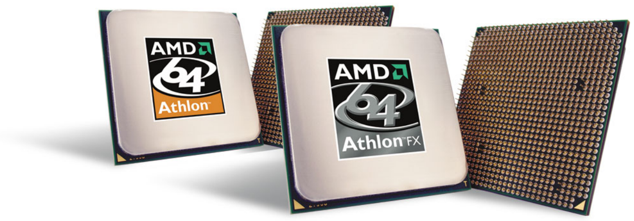

This doesn't mean that socket 940 or 754 Athlon 64 users need to start throwing out their gear. These new processors just provide a bit of fine polish to an already superb product and make it easier for OEMs and end users alike to follow their convergence models. On the FX-53 side, yes, there's a performance gain to be had over the older one due to the faster memory you can use, but when you're at the top, you can't expect to remain there for long anyway when it comes to CPUs. It's such a negligible difference that you needn't concern yourself with it.

Nothing has really changed in that the Athlon 64, in particular, the very expensive FX-53, is top-of-the-pile in terms of gaming. The 3800+ introduction isn't exactly cheap, either, and doesn't offer much better value for money, so if you're going to spend over $700, then you know which one to go for. The FX-53 processor is still below the price of the slower Pentium 4 Extreme Edition. Motherboard OEMs are already trickling in with VIA and NVIDIA-based socket 939 boards to help speed up the process. ▲

by Alex "Sharky" Ross

Specs: 0.13-micron, 128-bit memory interface and 512KB L2 cache

Specs (FX): 0.13-micron, 128-bit memory interface and 1MB L2 cache

CPU Ranking: 0 = Absolutely Worthless CPU CPU C 2.5 = Absolutely Average CPU CPU CPU CPU CPU 5 = Absolutely Perfect

FIC Condor

Condor
$300
FIC
(510) 252-777
www.fica.com

At first, we weren't overly impressed with the Condor as a next-gen SFF. For starters, the barebones box runs an Intel-manufactured, 865G/ICH5-based motherboard with zero overclocking or system monitoring features. Second, while the Condor sports one AGP 8X and two PCI slots, we had to wrestle with awkward angles and cramped cabling before we could seat our NVIDIA 5950 card. Third, the first 915 chipset-based SFF designs should be out soon, and Intel Extreme Graphics 3 will leave Springdale's IEG2 in the dust.

But once we cleared our minds of enthusiast preconceptions, we came increasingly to admire FIC's innovations. Styled with silver and charcoal plastic panels, we found the unit better looking in person than in print. On top of the SFF is a plastic plate, and set into it is a lever. Pull this, and the front, top, and left chassis panels—along with everything mounted to them internally—swing out and down on two hinges. These hinges are secured with buckles you can unclamp to separate the two chassis halves, although this will mean scrapping FIC's careful cable management. This is far more convenient than the typical XPC style of popping off the cover and removing the drive cage and cooling system.

The Condor has one external 5.25-inch and two 3.5-inch bays. FIC uses a sliding rail system for toolless installation, but this required a lot of shoving and jimmying. With only three drives and no OCing, the 200W power supply should be fine for most users.

There is a large grille on the floor of the chassis—which is not under the motherboard, right?—for cool air intake. The side-blowing TaiSol CPU heatsink fan draws this cool air through the low-profile northbridge heatsink, through the CPU heatsink, and out through the large mesh vent in the top of the chassis. According to FIC, total noise output is a mere 26dB.

The Condor's mobo is made by Intel, so you're trading speed for stability. We wouldn't take the Condor to a LAN party, but we'd be proud to plant several throughout the house for mainstream use. ▲

Specs: 865G/ICH5 chipset; Prescott-ready; bays: 1 x ODD, 1 x FDD, 1 x HDD; 10/100 LAN; 1 x AGP; 2 x PCI; front: 2 x USB, headphone, SPDIF; no FireWire

by William Van Winkle

Zalman Reserator 1

Reserator 1
$280
Zalman
(714) 530-0700
www.zalmanusa.com

We couldn't have envisioned a better, easier watercooling system than Corsair's HydroCool, but leave it to Zalman to take the imaginative leap. The Reserator 1 is a stunning, blue anodized aluminum tank standing nearly 2 feet tall. A small, silent pump sits in the base of the 2.7-quart chamber, but most heat transfer is enabled through convection. The total silence of the Reserator alone was enough to win our hearts.

But Zalman goes further by using its ZM-WB2 Gold water block, a copper heatsink with gold plating to help it retard galvanic corrosion. As such, while Zalman recommends distilled water, you can use tap water for coolant, too. Zalman includes a long stretch of soft silicone tubing and all the mounting gear you need for Sockets A, 940, or 478.

The instructions are adequate, and a newbie can expect to complete installation in 30 to 40 minutes. You should be able to do a second installation in less than 10 minutes. Speaking of which, Zalman throws in four tube clamps to choke off liquid flow in case you need to unscrew the connections and uninstall the Reserator. With most other watercoolers, you need to bring towels and slice through tubes to remove the equipment.

To test the Reserator 1, we replicated the tests we used for the HydroCool. With no big motor or fan, Zalman meets or beats Corsair even with the 3GHz CPU overclocked by 10% and working under a 95% burn-in load: the Zalman came in at 85 F; the Corsair shot to 89 F. For comparison, Intel's reference heatsink hit 120 F. The Zalman has an idle temp of 77 F.

Another perk is the inline flow indicator, a chamber with a cap floating in it. If the cap is vibrating, your coolant is flowing. With other watercoolers, you have to trust it's working. Of course, many other coolers include temp probes and readouts. Zalman offers neither, but for slight-to-moderate overclockers, that's fine. Extreme OCers can install a bay-mounted probe device.

Our only complaint about the design is that, unlike Corsair, Zalman does not include a backplane slot plate. Rather, the two round tube connectors mount into the slot with nothing but air around them.

The Reserator 1 is an artistic and technical masterpiece we can't recommend enough. ▲

Specs: 2.7-quart capacity; 23.3 inches high x 5.9 inches wide x 5.9 inches deep; 14.3 pounds; 5W pump; 300 l/hr flow; 1.5-foot max pump height

by William Van Winkle

CPU Ranking: 0 = Absolutely Worthless 2.5 = Absolutely Average 5 = Absolutely Perfect

ATI HDTV Wonder

HDTV Wonder
$199
ATI
(905) 882-2600
www.ati.com

ATI's new HDTV Wonder addresses demand for reception of free, OTA (over-the-air) broadcasts of digital TV signals on the PC, previously satisfied only by smaller companies. The HDTV Wonder's advantage is in its distinct software package that employs the best features of ATI's Multimedia Center and introduces a few unique touches. For example, the DTV player retains the same control interface as other Multimedia Center applications (including use with the bundled Remote Wonder), adding support for digital television, recording, scheduling, and pausing live DTV. Setup is straightforward, as ATI automates the entire process; the only obvious omission is an EPG (electronic programming guide) that displays programming data days in advance, a side effect of the medium's OTA nature.

Rather than recording transport streams, the HDTV Wonder can export directly to MPEG-2 or MPEG-4. ATI's Thruview technology facilitates levels of transparency, if you'd rather work on a document while watching HDTV in the background. And the board's NXT2004 demodulator is notorious for its signal sensitivity, ensuring the best possible reception. The card even comes with a UHF antenna ideal for receiving DTV broadcasts.

In addition to its digital tuning capabilities, the HDTV Wonder also receives analog programming through a separate antenna or cable connection. ATI's software suite can similarly record and time-shift analog television, capturing to MPEG-4, while the bundled Gemstar Guide+ EPG schedules recordings.

The HDTV Wonder's system requirements are minimal. You'll need a 1.3GHz Pentium 4 and at least 256MB of RAM. ATI recommends a Radeon 9500 or higher for optimal performance, and the Guide+ EPG requires an Internet connection for analog programming. Finally, visit www.antennaweb.org to verify that you're within range of an antenna broadcasting DTV signals. Most areas should get at least a few channels. Bakersfield, Calif., unfortunately, does not. ▲

Specs: DTV Tuning, analog TV tuning, and ATI NXT2004 demodulator; antenna and remote control included

by Chris Angelini

eVGA Personal Cinema FX 5700

Personal Cinema FX 5700
$250
eVGA
(714) 528-4500
www.evga.com

When it first debuted, NVIDIA's Personal Cinema lineup didn't compare favorably to ATI's All-In-Wonder family. Incremental enhancements have slowly evened the playing field, however, culminating in the Personal Cinema FX 5700, the flagship in NVIDIA's modest multimedia armada.

The card's ForceWare Multimedia software package is the most notable improvement, featuring a unified interface for watching TV, FM tuning, retrieving EPG (electronic programming guide) data, DVD playback, and accessing archived video. It even plays back transport stream files, the product of recording HD video from an HDTV tuner card. Ulead's DVD MovieFactory 2.5 SE and VideoStudio 7.0 SE are the only third-party applications bundled this time around, and those are used solely for authoring and editing video content.

Centered on NVIDIA's GeForce FX 5700 non-Ultra processor and 128MB of DDR memory on a 128-bit bus, the latest Personal Cinema won't blow through demanding first-person shooters. Far Cry, set at 1,024 x 768 with maximized details, runs at 26fps. UT2004, similarly configured (without antialiasing or anisotropic filtering), peaks at 54fps. Older games, such as Quake III, run much more smoothly, but it should be clear that the Personal Cinema FX 5700 is an average 3D performer.

It does excel in external connectivity, though. The back of the single-slot card features an input for cable or an antenna, an FM antenna input, a DVI output, a loop-back connector for routing audio to a sound card, and an interface for connecting the included breakout box. The breakout box further expounds on the card's options, providing composite video input, an RCA audio input, an S-Video input, an S-Video output, and component output for use with a high-definition television. An RF remote contributes to the product's home theater allure, even if it's reminiscent of ATI's old Remote Wonder.

Priced at $250 online, the Personal Cinema FX 5700 is forced to contend with ATI's All-In-Wonder 9600 XT, an arguably more refined card. Although it is unquestionably the best Personal Cinema effort to date, the FX 5700 falls just shy of its competition. ▲

Specs: NVIDIA GeForce FX 5700 GPU, 128MB of DDR memory, 400MHz RAMDAC, DVI output

by Chris Angelini

CPU Ranking: 0 = Absolutely Worthless 2.5 = Absolutely Average 5 = Absolutely Perfect

EPoX EP-8KDA3+ vs. MSI K8N Neo2 Platinum

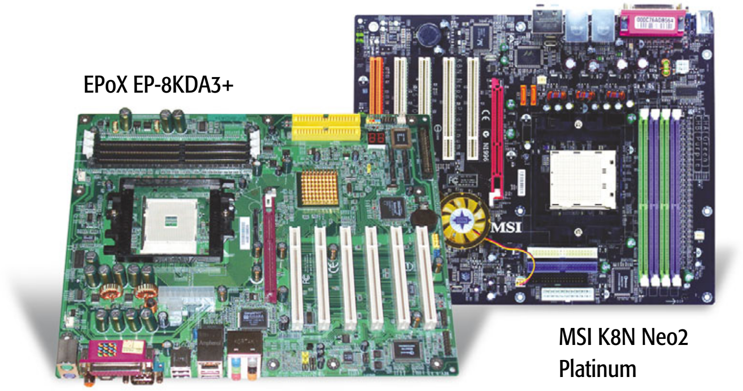
EPoX EP-8KDA3+

MSI K8N Neo2 Platinum

While AMD is making the awkward jump from 754- to 939-pin Athlon 64 chips, NVIDIA is working to follow suit with the nForce3 250Gb and nForce3 Ultra chipsets, respectively. Yes, the 939-pin format allows for dual-channel memory architecture and jumps from an 800MHz HyperTransport bus to 1GHz, but does this translate into a real world advantage?

We wanted to test a representative from each side of this bridge to look for similarities or differences while taking NVIDIA's scalding hot Gigabit Ethernet controller for a spin. Our test platform included an Athlon 64 3200+ (754-pin, 2GHz) and 3800+ (939-pin, 2.4 GHz), NVIDIA 5950 reference card, 250GB Hitachi Deskstar, and Windows XP SP1. NVIDIA supplied a dual Xeon server with IXIA's IxChariot traffic analysis app and an Alacritech Gigabit Accelerator specifically to flood the GbE connection to capacity.

EPoX EP-8KDA3+

The bad news here is about feature placement. The main and 12V power connectors are located right below the CPU, so you have to fight to keep your wires out of the fan, and the SATA 1/2 ports are jimmied between the AGP slot and CPU heatsink retention bracket, making them potentially hard to reach. Moreover, there is no FireWire support.

The good news is that performance was exactly where we expected. Earlier Web reviews showed this board with slow memory performance, but the KDA34514 BIOS that EPoX sent us right before our review seems to have remedied this. Interestingly, we did notice occasional drops in individual pair throughput rates that we didn't observe with MSI. This likely accounts for the slightly lower EPoX scores under Chariot.

Still, the EP-8KDA3+ gave us no serious reason to complain, and on our bench that's saying a lot.

MSI K8N Neo2 Platinum

Plenty of perks wait in this gem. You get NVIDIA's GbE connection plus a second GbE controller from Marvell. Realtek's 7.1 chip feeds to 1/8-inch analog jacks, as well as optical and coax SPDIF-Out. Like EPoX, MSI plants the SATA 1/2 ports behind the AGP card, but you have a second set of RAID-ready SATA ports elsewhere. The backplane has one FireWire and four USB ports, but there are headers on the board for another four of each connector.

The Phoenix BIOS provides for plenty of OC tweaks from voltages to clock ratios to HyperTransport multiplier. However, we wanted to experiment with MSI's CoreCell chip, which impressed so well in our Intel 865 board roundup. Out of the box, CoreCell is disabled, but you can set it from Private (1% OC) to Commander (11%). The top stable setting we could get from CoreCell, however, was Captain (5%) with an 800MHz HyperTransport and only Captain (3%) at 1GHz. We'll stick with manual OC settings, thanks.

Overall, performance is good, we had no technical problems, and MSI's feature set is nearly unbeatable. Subsequent nForce3 Ultra boards will be hard-pressed to top this act.

The Verdict: Great Gigabit, Modest Memory

As our PCMark tests show (subscribers can see the chart at www.cpumag.com/cpuaug04/benchmarks), when all other factors are equal, dual channel screams past single channel, but you need an application that stresses RAM in order to see it. Similarly, if the bump in HyperTransport from 4X to 5X is a big deal, we have yet to see why. Thus ends the 939 vs. 754 discussion.

The 250Gb tests prove the PCI bus is no longer a valid networking option for high-speed use. To compare, we planted SMC's SMC9452TX 32-bit PCI GbE adapter on the EPoX board and realized only 657.531Mbps average throughput. Broadcom's integrated chip shows similar sloth. ▲

by William Van Winkle

K8N Neo2 Platinum
$169.99
MSI
(626) 913-0828
www.msicomputer.com

EP-8KDA3+
$165
EPoX
(714) 680-0898
www.epox.com/USA

Specs: nForce3 250Gb, Socket 754, 6 x PCI, AGP 8X, 3 x PC3200 slots, 2 x PATA connectors, 6 x SATA connectors, Sil3114 SATA RAID, ALC850 7.1 audio

Specs: nForce3 Ultra, Socket 939, 5 x PCO, AGP 8X, 1 x FireWire, 4 x PC3200 dual memory slots, 2 x PATA connectors, 4 x SATA connectors, SATA RAID, ALC850 7.1 audio

CPU Ranking: 0 = Absolutely Worthless 2.5 = Absolutely Average 5 = Absolutely Perfect

Logitech MX510 Performance Optical Mouse

MX510 Performance Optical Mouse
$49.95
Logitech
(800) 231-7717
(510) 795-8500
www.logitech.com
CPU CPU CPU C

Logitech has a knack for creating mice that are both stylish and highly functional. The MX510 Performance Optical Mouse is no exception. Although the colorful and shimmery body of the mouse is sure to attract attention, the real value in this mouse is its optical technology, eight programmable buttons, and ergonomic design.

The MX510 is loaded with eight buttons: left, right, scroll, Cruise Up and Cruise Down (for scrolling quickly at the touch of a button), Quick Switch (for toggling among open windows), and Back and Forward buttons (for navigating either Internet Explorer or Windows Explorer). The bundled MouseWare utility lets you reassign each of the MX510's eight buttons with one of 33 functions, such as Close Application, Copy, Paste, and Undo.

This mouse is most comfortable if you have large hands, but those with medium-sized hands will still find it a comfortable fit. Left-handed users are left out in the cold with this mouse, however. The position of the Back and Forward buttons, as well as the sloped and sculpted body of the mouse, accommodates only right-handed users.

The left side of the mouse is sculpted to support the thumb, while the right side of the mouse is slightly curved, providing a comfortable groove for the right pinky finger. Even the top of the mouse slants slightly to the right to provide the most comfortable angle at which to hold your wrist.

The optical performance of the MX510 is impressive. In my tests the optical sensor was highly responsive on all surfaces on which I used the mouse. Only the most reflective surfaces affected the accuracy of the on-screen pointer.

The MX510 comes in either an iridescent red or blue. This mouse isn't cheap, but it does offer a very comfortable user experience and outperforms most similar optical mice. ▲

by Kylee Dickey

Specs: Eight buttons, ergonomic design, highly responsive optical sensor

Cooler Master CM Stacker STC-T01

CM Stacker STC-T01
$199
Cooler Master
(909) 673-9880
www.coolermaster.com
CPU CPU CPU CPU

You may well keep this tower until your dying day . . . or at least until the BTX form factor vanishes. The CM Stacker STC-T01 aluminum beauty features 11 5.25-inch bays, all external and covered by hollow bay plates covered in mesh. Why mesh? Glad you asked. Apart from helping to vent the entire front panel, the open design provides a cool air source for the 4-in-3 Device Module. This is a cage that occupies three of the 5.25-inch bays yet holds four hard drives and has a 120mm cooling intake fan mounted to the front. This is in addition to an 80mm exhaust fan in the top of the case and another 120mm exhaust fan at the rear. Moreover, there is a long vent grille spanning most of the chassis' floor and punch-outs for two 80mm fans below the backplane slots. Alternatively, you can mount your power supply in this bottom position near the case floor and put your two extra 80mm fans up where the PSU normally goes.

The left panel is dominated by a circular vent with a removable dust filter behind which you can mount another case fan that would rest above your CPU. Even with the optional 13-inch cross-flow fan mounted between your drives and the motherboard, Cooler Master does a remarkable job of keeping this tower hushed considering that you're starting with at least four or five fans.

Once you screw a pair of plastic rails onto each drive, you merely flip open the latches on the side of each bay, slide in the drive, and close the latches. The single FireWire, two audio, and six USB ports are thankfully located at the top of the front panel. Immediately below these is a bay plate cut for a floppy drive, but no 3.5-inch to 5.25-inch adapter kit was included.

This brings up the tower's only shortcoming (besides the lack of a dust filter on the bottom grille): In all that real estate, you only have four hard drive bays. Cooler Master no doubt hopes you'll opt for the various thermal sensors, regulators, VU meters, and such that it sells to fill these bays once you've finished with two or three optical drives, but there should at least be adapters included for additional drives.

There's a lot of screw work with this chassis, but between rock-solid construction; coaster wheels; and support for ATX, E-ATX, and BTX form factors, Cooler Master has yet another smash hit with the CM Stacker. ▲

by William Van Winkle

Specs: 11 5.25-inch bays; ATX-, E-ATX-, BTX-compatible; support for eight fans; three fans (2 x 120 mm, 1 x 80 mm) included; front ports: one FireWire, mic, headphone, six USB ports

CPU Ranking: 0 = Absolutely Worthless CPU CPU C 2.5 = Absolutely Average CPU CPU CPU CPU CPU 5 = Absolutely Perfect

Anand's Corner

by Anand Lal Shimpi

Intel's New 925X & 915

Anand Lal Shimpi has turned a fledgling personal page on GeoCities.com into one of the world's most visited and trusted PC hardware sites. Anand started his site in 1997 at just 14 years old and has since been featured in USA Today, *CBS' 48 Hours and* Fortune. *His site—www.anandtech.com—receives more than 55 million page views and is read by more than 2 million readers per month.*

Intel has been talking about, and I've been leaking information about, its next-generation PCI Express chipsets for months now. Originally code-named Alderwood and Grantsdale, the official names for the new chipsets are 925X and 915, respectively. The specs of the chipsets are fairly widely known by now, but here's a brief overview for those of you that have better things to occupy your time.

Both chipsets are basically identical in their support for the following features:

- 1 – PCI Express x16 port
- 4 – PCI Express x1 slots
- 8 – USB 2.0 ports
- 4 – SATA ports
- 1 – PATA port
- Dual Channel DDR2-533 Memory Interface

There are a couple of differentiating features: [1] 925X only supports DDR2, while 915 has support for both DDR and DDR2; [2] 925X has some "performance enhancements" that reduce memory access latency; and [3] 915 is available with or without integrated graphics (915G or 915P); 925X is only available without.

The new chipset does away with AGP support in favor of the new PCI Express x16 interface for graphics. From a performance standpoint, I wouldn't expect anything big out of PCI Express to begin with—the added bandwidth that the interface offers won't find itself being used today. How quickly the GPU guys and the game developers end up taking advantage of the added bandwidth of PCI Express has yet to be seen, but hopefully the lack of usage of AGP 8X's bandwidth over 4X will not be indicative of things to come.

> **The new chipset does away with AGP support . . .**

Other than PCI Express, the next big thing to talk about with these chipsets is their support for DDR2 memory. Basically, DDR2 memory allows for lower core clock speed operation of the actual memory devices through a doubling of the amount of data that's prefetched every cycle. The benefits are higher potential memory speeds, which is even evident today as DDR2 is making its debut on 925X/915 at an effective 533MHz clock. Unfortunately, the initial DDR2 memories available on the market are running at much higher latencies than the fastest DDR modules. With the added bandwidth offered by DDR2 not being of much use to us today (the P4 is still FSB/clock speed limited), the higher latency of DDR2 means that a 925X system with DDR2-533 performs about the same as a current generation 875P system with DDR400 memory. DDR2 pricing is also not something to get too excited about, but that is to be expected since the memory has just recently hit the market.

The performance enhancements that 925X offers over 915 are similar to the differences between the 875P and 865PE chipsets—again, not amounting to much in the real world.

Intel has significantly improved its integrated graphics core by quadrupling the number of pixel shader pipelines. While a quadrupling of anything sounds impressive, considering that its old integrated graphics solution had a single pixel pipe, going to 4 isn't really impressive, it's just about time. Clock speeds have also improved performance, and integrated graphics is one area where the added bandwidth of DDR2 will show performance improvements today. It's just that integrated graphics platforms are probably the least likely to use expensive DDR2-533 memory, thus so much for that present-day advantage.

With the new chipsets Intel is also introducing its new LGA-775 CPU socket. The new CPU interface is actually causing quite a bit of controversy as the socket is much more prone to damage than the previous sockets. Unlike current Socket-478 interfaces, the pins for LGA-775 aren't located on the CPU itself. No that's flat, the pins are instead located on the motherboard and (with 775 of them to count) bending them isn't too difficult.

The new 925X/915 boards will also be featuring an ATX-Extended power connector (2x12 pins vs. 2x10). The added power pins are to deliver the necessary power to the PCI Express x16 slot, which can provide more power to graphics cards than regular AGP slots could. At this point in time, it seems that only NVIDIA's NV45 needs the added power, but that'll probably change with the next generation of GPUs.

Technologically, the new chipsets are definitely interesting. However, things could be better with regard to well-performing upgrades for today. ■

Talk back to Anand@cpumag.com.

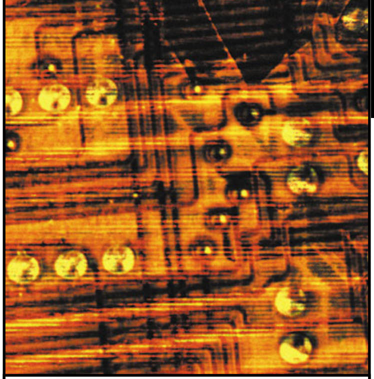

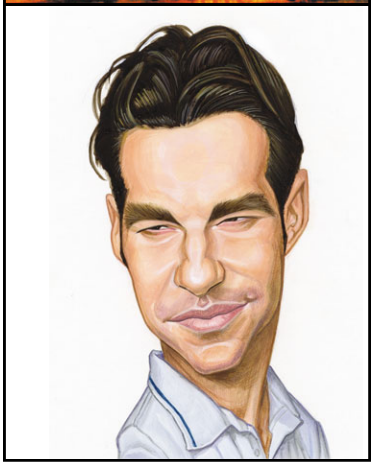

Disrupting Reuters' newswire with a cheery Christmas greeting at age six, Alex "Sharky" Ross became an avid computer user/abuser, eventually founding popular hardware testing/review Web site SharkyExtreme.com. Exposing shoddy manufacturing practices and rubbish-spouting marketing weasels while championing innovative products, illuminating new technology, and pioneering real-world testing methods was just a front for playing with the best toys. The site acquired, he left in 2001. A London native and London School of Economics graduate, Alex currently swims in Silicon Valley.

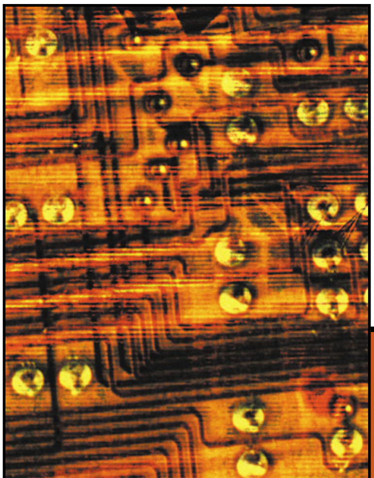

The Shark Tank

by Alex "Sharky" Ross

MXM: Mobile GPU Swap?

One of the biggest factors that usually keeps me from buying a notebook is that once you drop that $1.5-2K on one, you end up being stuck with it. In other words, there's no room for upgrading other than memory and hard drive, neither of which is going to dramatically improve performance in six to 12 months. Being a gamer, I do enjoy the occasional blast/drive with some DX8/DX9/Open GL titles. Hence, I'm not too keen on being stuck with a previous-generation 3D graphics chip on the desktop, so why should it be any different on a notebook?

Indeed, some companies, such as Alienware and Dell, have been taking steps to future-proof their notebooks by introducing upgradeable graphics capabilities to allow users to upgrade from ATI M10 to M11 graphics, for example. The trouble is that there is no consistent standard such as an AGP port as there is on the desktop, so OEMs end up implementing their own interfaces, which are not compatible. With PCI Express just around the corner (seems to be one of those never-ending corners, I know), NVIDIA is readying its MXM (Mobile PCI Express Module) in an effort to bring a new standard to the notebook world.

End users may have the choice to switch from an ATI to an NVIDIA chip or vice versa.

ODMs (original design manufacturers) such as Quanta, Wistron, AOpen, and FIC are already facing design challenges and changes by going from AGP to PCI Express, so what's another NVIDIA spanner in the works? During a recent meeting with NVIDIA at the company's San Jose HQ, representatives claimed that the aforementioned ODMs couldn't be happier with the supposed benefits that MXM will bring. First, using a standard allows for an easier upgrade path with swappable graphics. Second, a standard would allow the same notebook design to be applied to many different products/SKUs. Third, a standard could help reduce the time it takes to get to market with newer products. And finally, MXM is compatible with graphics from "any vendor," as NVIDIA puts it. In other words, it doesn't matter whether, as a notebook manufacturer, you decide upon ATI or NVIDIA graphics, as either will work. This makes for an interesting situation with the current back-and-forth shenanigans between the performance of ATI and NVIDIA hardware. It means that end users may have the choice to switch from an ATI to an NVIDIA chip or vice versa depending upon whose technology is better at the time.

Motherboards that use an MXM slot (with 230 pins) will sport a small 2.9-inch-long x 0.2-inch-wide x 0.44-inch-deep connector. The width will vary from 0.2 inches wide to as small as .06 inches, depending upon the type of application desired. The MXM slot will feature full 16-lane PCI Express signals and is able to run the display I/O signals running through it. There will be three choices of MXM modules: MXM-I, MXM-II, and MXM-III to suit different SKUs. MXM-I is the smallest of the three at 2.8 inches wide x 2.72 inches long with a power budget up to 18 watts. With a maximum allocation of 1.4 inches x 1.4 inches for a GPU core and up to 128MB, it is designed to be used in the thin and ultralight sector. MXM-II is slightly larger at 2.92 inches x 3.12 inches with a greater power budget (up to 25 watts) and has support for more memory at 256MB. It will likely be the mainstream part. For the desktop replacement sector, MXM-III makes the most sense with its 3.28-inch x 4-inch size and 35-watt capability. It also supports up to 256MB of memory, but more importantly, has enough room for a 1.6-inch x 1.6-inch GPU.

The advantages to ODMs by negating the extra design time for those proprietary modules and the improved and easier upgrade paths are clear, but here's a question: Will this open standard from NVIDIA receive support from across the industry? Will ATI and S3, for example, produce MXM graphics modules? There's been no official statement from either of those players, but should the Taiwanese ODMs choose to, they could certainly go ahead and design their own MXM cards that utilize either ATI or S3 hardware. The old days of multimillion-dollar design wins might turn out to be slightly less stringent if end users are able to choose what graphics they want for themselves. ■

Swap mail at sharky@cpumag.com.

by Kyle Bennett

WoOD$+OcK PhoR G33k5

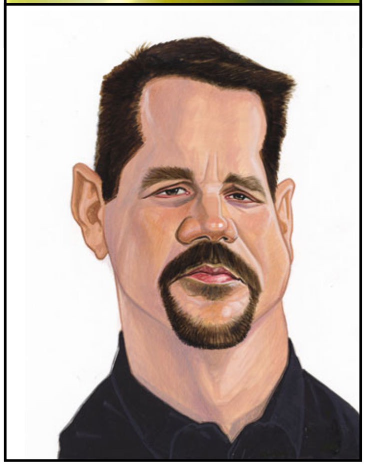

Kyle Bennett is editor-in-chief of HardOCP.com (hardocp.com), one of the largest and most outspoken PC-enthusiast sites on the Web. HardOCP.com is geared toward users with a passion for PCs and those who want to get cutting-edge performance from their systems. Beware, though, Kyle is known for his strong opinions and stating them in a no-nonsense manner while delivering some of the most in-depth reviews and PC hardware news on the 'Net.

It's getting to be that time of year again, and there's nothing like August in Texas: 100 degrees outside, temperatures baking, clear skies for miles, longhorns standing in the shade, spittin' "chaw" on the sizzling sidewalk. And, oh yeah, more than 3,000 gamers packed into a luxurious hotel fragging their asses off.

QuakeCon is a long-standing gamer tradition and industry icon among LAN parties. While some will argue that QuakeCon is a LAN-gaming mecca that happens to have a huge party attached, others will suggest that it's a big party with a LAN that you can attend while you're there. Trust me, we have been doing this for years, and the truth lies somewhere in the middle. There's more than enough of each to go around, and at the very least, when/if you get tired of one, you can go back to the other.

So is QuakeCon all about the gamers in the BYOC (Bring Your Own Computer) area and the hardcore that are there for tournament dollars and prizes? Hardly. QuakeCon is about community and the folks that bind that community together getting out from behind their keyboards and spending some face time with each other. Certainly you will see lots of old friends and meet lots of new friends. Last year, somewhere from 3,000 to 4,000 people, other than those gaming in the BYOC, showed up to share in the festivities.

id Software is the driving force behind QuakeCon, but even it takes a back seat to the heart and soul of the operation. You see, QuakeCon is run by volunteers. That's right, just a bunch of folks that want to be involved in helping give back to their community. If you have never seen a few tables full of gamers cranking out thousands of CAT 5 cables, QuakeCon is where you can see it. The person that greets you at the registration table, the guy that assigns a sticker to your equipment, and even the guy that kicks you in the shin for line-jumping are all there because they want to be, not because they are getting paid. The volunteers bring the sweat off their brow, and id Software brings its checkbook and a library of computer games that keeps everyone very busy well into the night.

And because the event is put on by volunteers, it makes perfect sense that you shouldn't be charged for just walking through the door. And that's just the way it has been done since the event's inception. The sponsors won't pay for your hotel room, but they will give you a place to game, a place to relax, a place to hang out, and about as much swag as you can pack into your car.

Outside of id Software and the volunteers, however, the sponsors are the ones that make it happen, and we thank them for that. It's not the hard sell that you get other places, though. There's a whole lot of "try it before you buy it" going on. QuakeCon is a place to learn about new technology in the market instead of having someone trying to sell it to you. QuakeCon is a place to share ideas with tech giants in the industry while they share the worthiness of their wares. Quite frankly, QuakeCon seems to be as much of a learning experience for the sponsors as the attendees.

> **Quite frankly, QuakeCon seems to be as much of a learning experience for the sponsors as the attendees.**

So pack your bags and get yourself to Dallas for QuakeCon (www.quakecon.org) this year Aug. 12-15. Did I mention the area is going to be a little bigger this year, as well? Over 175,000 square feet of floor space, and that doesn't even count the acres of river walk in the air conditioning under the glass ceiling!

And as has become tradition, the [H]ard|OCP Hardware Workshop will take place on Saturday afternoon. It's always a fun mix of hardware information and tons of giveaways. We should have room for well over a thousand folks this year. And who knows, we might be able to tell you exactly what you will need to build that DOOM3 box that you have been dreaming of for a couple of years now. ■

You can talk with Kyle at kyle@cpumag.com.

Lighten Up.

It wasn't so long ago that if you wanted to bring your data with you, your hardware had to come too. Well, that's all in the past now thanks to the new SanDisk Cruzer™ Mini. At under half an ounce, the Cruzer Mini goes anywhere, and with up to 256MB of flash memory, it takes a lot with it. Hi-speed USB 2.0 compliant, SanDisk's Cruzer Mini lets you transfer data faster than ever before. And at less than 8mm thick, all it needs is a single USB slot. Just drag, drop, and go. It's all about convenience – no extra hardware, no worries.

So whether you're installing utilities or updating print drivers for the umpteenth time, you can take your files with you wherever you go – and you won't even need your dolly.

For more information on SanDisk Cruzer Mini, visit www.sandisk.com/smart

HI-SPEED CERTIFIED USB
HI-SPEED USB 2.0

SanDisk®

SanDisk and the SanDisk logo are registered trademarks and Cruzer is a trademark of SanDisk Corporation. Other brands and products are trademarks of their respective holders. ©2004 SanDisk Corporation. All rights reserved.

PC Modder

Tips & Tutorials

Modding does the body good. A PC's body anyway, inside and out. Here you'll find hardware, firmware, tools, tips, and tutorials for modding your rig's performance and appearance. Send us your own mod-related tips and ideas at modding@cpumag.com.

Modding enthusiasts have a penchant for the latest toys. Fast processors, powerful video cards, silent SFF enclosures, and radically lit motherboards are all fair game when it comes to a modder's creative mind. Recognizing the appeal of unconventional customization, an entire industry has emerged to support the community.

Coolmax's CF-400B operates without a fan, requiring ventilation to exhaust excess heat.

Mods & Ends

"Silent" Power Supply Roundup

While many modded systems project flamboyance and extravagance, others purposely blend right into their environments, such as HTPCs. Exterior design and hardware selection are both important in a successful implementation, but acoustics determine if your friends will hear the movie or an array of high-speed 80mm fans buzzing away.

Sapphire sells several passively cooled graphics cards, and there are various sources for CPU coolers, but power supplies generally persist as perpetrators in the fight to reduce noise. Recent fanless units by Silverstone Technology and Coolmax change that. Zalman also offers what it purports as a noise-free solution, although it uses one thermally controlled fan.

Coolmax's CF-400B is one of the most promising contenders. It has a sleeved 20-pin power header and six large 4-pin power connectors, a SATA connector, two smaller power connectors, and an auxiliary 4-pin power connector. Further, the CF-400B's aluminum chassis is riddled with rectangular holes to facilitate ventilation. Stagnant heat does cause the CF-400B to run fairly hot, and it's a solid bet temps inside your case will increase as a result. Moreover, at $125, the CF-400B is a costly upgrade.

The Silverstone SST-30NF tops out at 300W and sells for about $155. The unit is fanless but uses heatpipe technology with an aluminum case to transfer heat from the PSU core out to exterior fins. The unit comes with two SATA power connectors, six 4-pin power headers, two smaller headers, a sleeved 20-pin ATX power connector, and one 4-pin auxiliary connector. Silverstone products generally boast a higher degree of polish. The SST-30NF's silver chassis is no exception, and the output specs reflect similar quality. Though only rated for 300W, the unit pushes 18A on the +12V rail.

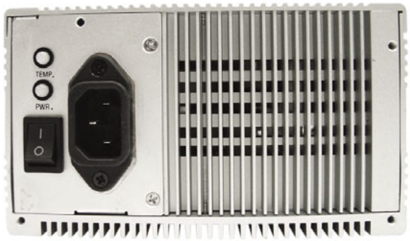

The fanless SST-30NF from Silverstone offers plenty of bling for the buck.

At less than $100, Zalman's ZM400B-APS is the least expensive option. It's also the only one with a fan. Rather than claim complete silence, Zalman uses a thermal sensor to monitor the PSU's temp and dynamically control fan speed. The result is a more reliable operation but greater noise. In a home theater environment, though, the extra stability might be worth it. The unit is even more attractive considering its included fan connector, which supports case fans with 3-pin power connectors.

Though it uses a fan, Zalman's ZM400B-APS is nearly inaudible in most cases.

Fashionably Fresh Firmware

Sony DRU-530A and DRX-530UL 8X DVD±R/RW Drive

Sony recently released several new firmware updates. Version 2.1A, the most significant, adds 8X DVD-R burning on 8X media and improves read compatibility with recorded dual-layer discs.

sony.storagesupport.com/dvdrw
/firmware.htm

Nikon Coolpix 3700 3.2MP Digital Camera

Nikon's v.1.2 firmware corrects problems found while saving images to certain brands of SD cards. It also fixes the brightness level issue in movies recorded under certain lighting conditions.

www.nikonusa.com

Handspring Treo 600

Installing the v.1.13 firmware update improves data connectivity, the way Web pages render, and audio quality, plus adds echo cancellation for car use and enables a privacy mode for incoming SMS messages.

www.palmone.com/us/support
/downloads/treo600.html

by Chris Angelini

Building "Spike"

The Diary Of A Modder's Creation, Part II

Last month (July CPU*, page 31), we began following the progress of "Spike," a modded system that Gareth Powell is creating from scratch. Last month Gareth set the foundation for Spike by sketching rough drafts and planning the hardware to use. This month he begins the assembly. Gareth will unveil Spike later this summer at the Cyberathlete Professional League's opening ceremonies.*

No matter how hard I try, I cannot control everything.

With the design and development of Spike done, I had to face the gritty truth that this design has to become reality; the sketches and ideas in my head have to become part of a real physical PC. The excitement of the design phase is over, and the long process of building Spike has begun. Deadlines and working dates now become real, and stress can take a hold if I don't watch what I'm doing.

I thrive under pressure, and in some respects I do my best work with a deadline looming over my head. I get a kind of rush from knowing that I may not make a cut-off date. From experience, however, I know that not making a deadline is something I don't want to do. I now plan everything so that I have ample time for setbacks and unexpected delays. Yet, even with the best planning in the world, I can't control everything.

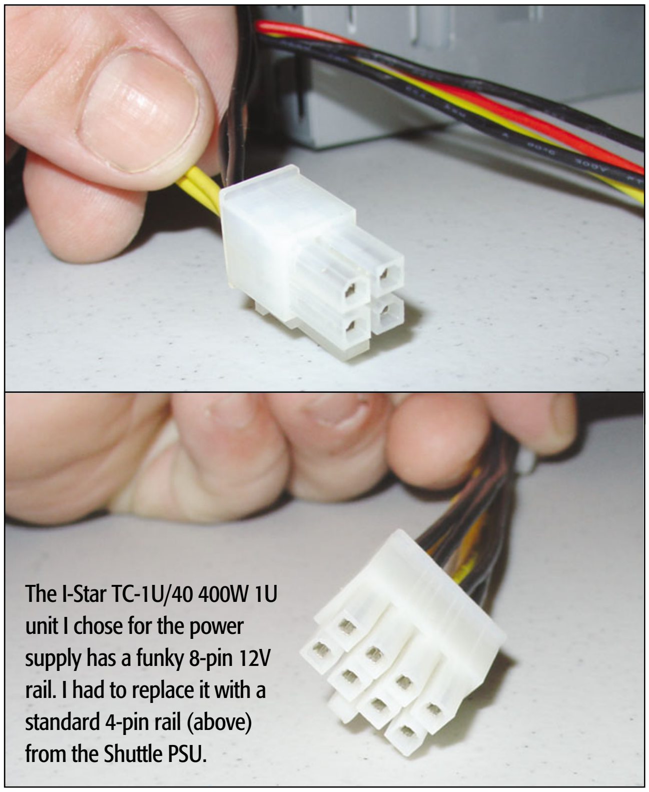
The I-Star TC-1U/40 400W 1U unit I chose for the power supply has a funky 8-pin 12V rail. I had to replace it with a standard 4-pin rail (above) from the Shuttle PSU.

I live in Burbank, Calif., the media capital of the world. Warner Brothers, Disney, and Universal Studios are all in the area. Therefore, it's a given that we have a concentration of special FX studios, set designers, and set builders, along with myriad supply stores and special manufacturing shops. This is perfect for a case modder because on any given day I can find a laser cutter, CNC mill, or water jet machine that can cut a special part for me. Plus, I can usually find whatever material I need to make my cases. But there is a downside to living here.

Rigid Foam Blues

Right now we are in the middle of the film-production season and the beginning of the new television-shooting schedule. Therefore, some materials can become hard to find, such as rigid sanding foam. Most film sets are made from rigid foam because it is easy to sculpt and can be painted to look like any type of surface. Rigid foam just so happens to also be what I need to make my side panels. But with all the major suppliers in the area on back order, I will have to wait. Economics dictate that the big production companies get precedence over my measly half-sheet order.

For some reason I also have had a problem getting the aluminum sheet I need. Because the spikes for this creation have to be cut and banged out by hand, I specifically need zero temper aluminum. This metal is untreated and soft, which makes it perfect for manipulating it by hand. During the process of pounding it into shape, it will harden and hold its shape. It also polishes up nicely. The problem is that this stuff isn't stocked in your local hardware store because it isn't normally used in construction. So I had to find a sheet at a metal supply store, of which there are plenty in the area. However, not one had it in stock. So here I am, at the start of a project that I can't get the foam for to work on the side panels or the metal I need to make the spikes.

With Murphy's Law in full effect, I trudged on in my search for metal and foam. I was looking for 063 aluminum, but after two days of hunting, the closest I could come up with was a 090 sheet. I bit the bullet and picked up the thicker metal. It will be a little harder to work with, but at this point, beggars can't be choosers when it's down to the wire.

Back To The Drawing Board

With the unexpected delay, I had some time to refine my design. I had my artist, Travis Jackson, do a few renders of the unit in scale so I could show others what I'm trying to accomplish. Spike will be a small but powerful unit when it's done, running a 3.2GHz Pentium 4 Extreme Edition (courtesy of Intel) on a Shuttle SB75G2 XPC motherboard. I have used a lot of Shuttle boards recently, and the small form factor is so much more versatile than the standard ATX board. Even though the XPC board doesn't equal ATX in stature, it matches up in power and performance, which makes it a good choice, especially when working with shapes that don't conform to the boxy case standard.

For video, I've chosen an NVIDIA GeForce FX 6800 card. This system will be a full-fledged gaming rig, so having anything less than the best would be pointless. The full-sized graphics card and its two-slot cooler require that I adjust the design a little to fit the card in the case. Therefore, I'm using an A-side AGP riser card, so the graphics card will sit parallel with the plane of the board, providing it with airflow for cooling. I typically use risers made by Adex Electronics. The risers are 8X AGP compatible and can be custom configured if needed.

The power supply that comes with the XPC is rated at 250W, but the draw of the 3.2GHz EE, GeForce FX 6800, hard drive, and lighting is too much for the standard XPC PSU. I've opted to use an I-Star TC-1U/40 400W 1U unit. This power supply has enough output, so the system won't have any instability due to lack of power. The power supply isn't exactly standard, though. It has a funky 8-pin 12V rail, so I have to replace it with the standard 4-pin from the Shuttle power supply to make it compatible with the board I'm using.

Spike will be 11 inches high x 12 inches wide x 21 inches deep when it's completed, which makes it an easy fit on anyone's desk. It is also constructed of aluminum and fiberglass, so it's ultra light and easily transportable to any LAN party.

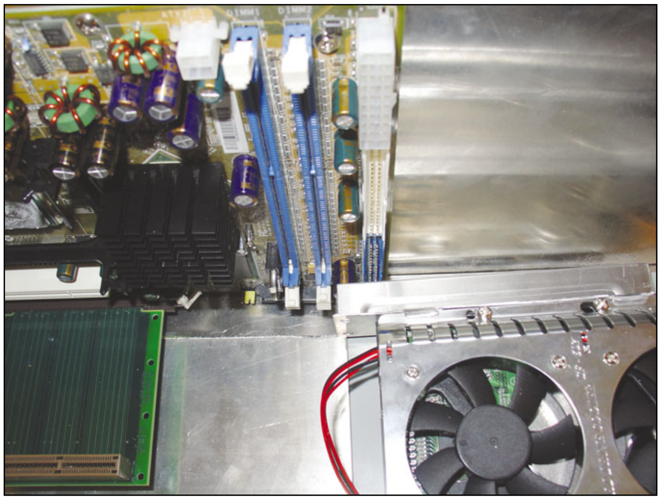

I've got problems. Due to an error in calculations, the CD-ROM was colliding with a capacitor on the motherboard.

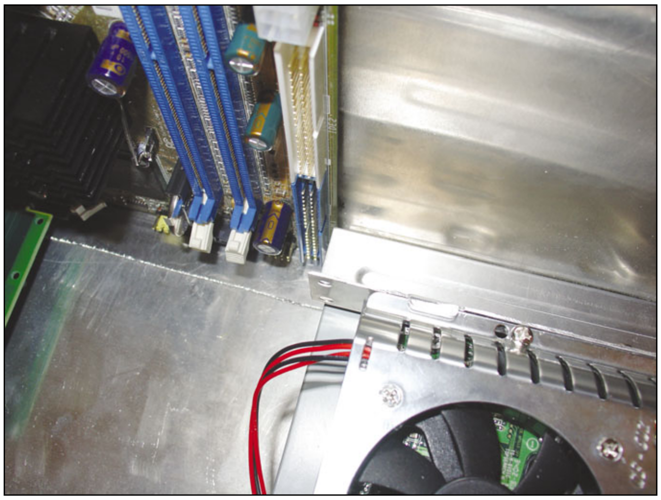

Problem solved. After some thought and much frustration, moving the board back helped make things more comfy.

Let The Light Shine

While I was waiting for the metal, I decided to get a head start on Spike's lighting. Standard cold-cathode lighting comes in basic colors, but as any artist knows, there are many more shades than colors. Blue, green, and white cold cathodes are really vibrant, and the red cathodes on the market are actually more pink than red. Therefore, I have taken on the task of making my own tinted cathodes using some industrial lighting gels. I found some white cathodes online that came with a sheath (like most cathode kits), and I visited a local industrial lighting store for a selection of red gels. By mixing and matching the gels, I was able to get the exact color I was looking for. I then cut the gels to size and taped them around the cathode sheaths. I was surprised how well this method worked, as I got the color I wanted and with relative ease.

After a few days, I finally got my metal and was able to start constructing the frame. For this I used 0.75-inch aluminum tubing bent to the correct radius. Aluminum is a soft metal, so bending it into the required radius without creasing the metal is a challenge. I enlisted the help of a friend who's an experienced metal former. For some reason the tubing kept creasing in a bender, so we painstakingly tapped out the tubing in a jig to the correct radius.

I learned how to weld in high school, and although I can weld steel, aluminum is a different matter. Aluminum welding is an art form in and of itself, taking years to master. Therefore, I had a friend weld the frame. I would have liked to do all the construction on this unit myself, but the time allotment for this project doesn't allow for me to learn aluminum welding and still have enough time to build the entire unit. Although the design is simple, the process of taking all the different elements of the enclosure and combining them into one enclosure is the real challenge.

The next step was filling out the frame, starting with the CD-ROM drive and hard drive bays, which in true modding fashion I butchered from the Shuttle case. I am using a Lite-On black bezel CD/DVD combo drive and a half-height 80GB Maxtor drive for Spike. To mount them I placed them upside down in the drive bays and fitted them to the case. The drives will be very visible in this case, and I didn't think that the Shuttle bay looked that good. I cut down the tray assembly into two brackets, which looked much better but still needed something to set it off. I hunted around in my parts closet and found a polished aluminum hard drive cooler. The cooler fit the style of the case, and it topped off the drive bays perfectly.

I mounted the cooler to the hard drive and flipped the drive so that the cooler would be the top element of the tray. It's perfect for this case, even though the drive is now mounted upside down. I don't think this will affect the performance, though. I have mounted drives in funky places before. My Technotone mod may have the funkiest drive mounting system in history. The slot-load DVD drive was mounted on a custom bracket under the LCD on the front so that you can load a disc through the opening. The logical place to put the hard drive would have been on top of the DVD drive, so I placed the drive on top of the unit at a 45-degree angle. Surprisingly, it didn't hurt the drive at all.

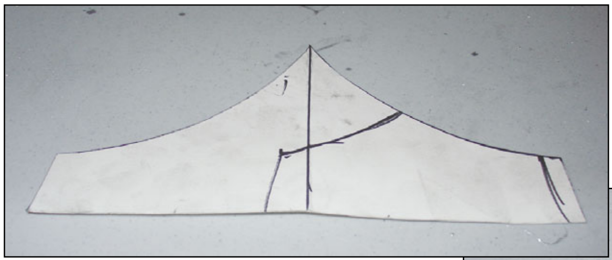

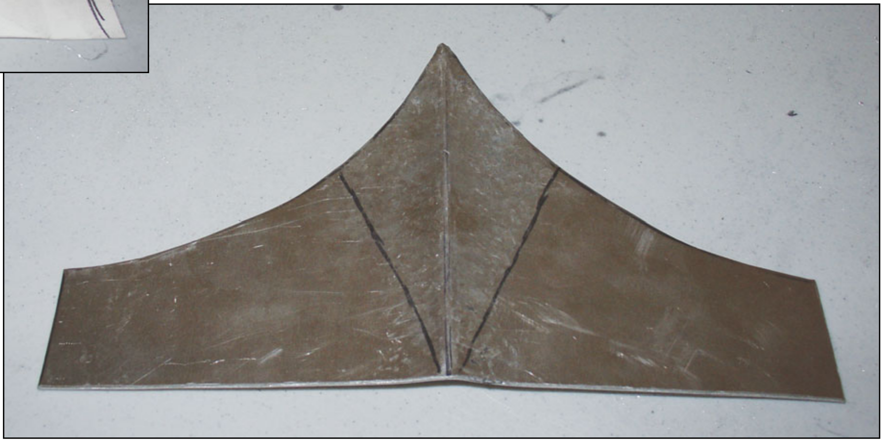

Here's a look at the template (above) I used to create a prototype spike for the case.

Panels To Mount

With Spike's drive in place, I had to make and mount the case's frame panels. This included the motherboard tray, backplane, and blocking panel. To do this, I made cardboard templates and fit them to the frame. The blocking panel was easiest. Its purpose is to fill in the internal box of the case. After fitting the template to the frame, I traced the shape of the panel onto commercial-grade 063 aluminum. Commercial grade is hardened, and because these pieces aren't being formed or bent, it makes sense to use the harder (and cheaper material). When tracing the pattern on the aluminum, I intentionally made the piece larger than it had to be so I could grind the metal down later to fit. I then riveted the blocking panel to the frame.

The motherboard panel was a little harder to place due to a cross member that will strengthen the case and serve as the mounting for the backplane. The motherboard tray also ties into the case's base for stability via a riveting lip. Again, I started with a cardboard template and transferred the shape to the aluminum sheet. I then cut the shape out. The lip needed to be bent, so I took it to a local shop and had it bent in a press brake. This gave the lip the bend I needed to fit it to the frame. After fitting the panel to the case, I had to drill and tap the motherboard mounting holes. I measured where I wanted the motherboard to be, marked the mounting holes with a punch, and drilled those suckers out. Now I could rivet the motherboard tray in place and grind it down to fit.

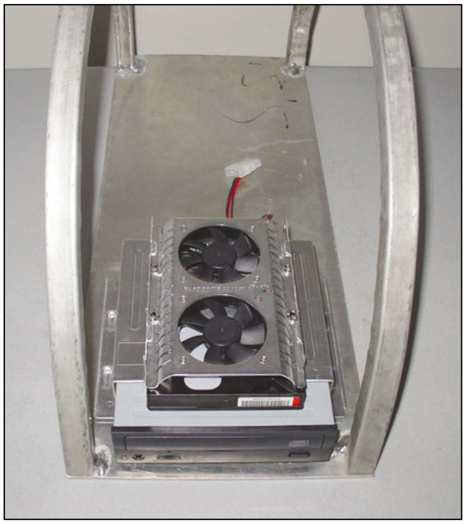

Fitting the CD-ROM drive in just the right spot.

I decided to give the motherboard a test fitting with the CD drive and the power supply. It was here that I found that I had made a critical error in my calculations. I hadn't accounted for the height of the capacitor on the board and left enough room for it. The CD-ROM was colliding with the capacitor and cutting off the IDE channel. I had two options: relocate the capacitor or move the motherboard.

If I relocated the capacitor, I couldn't just leave it floating around. I would need somewhere to place it. I also would still be unable to use the primary IDE channel. I decided to move the board back an inch, but that would also create another problem. The cross member at the rear of the frame isn't only there for stability; it's

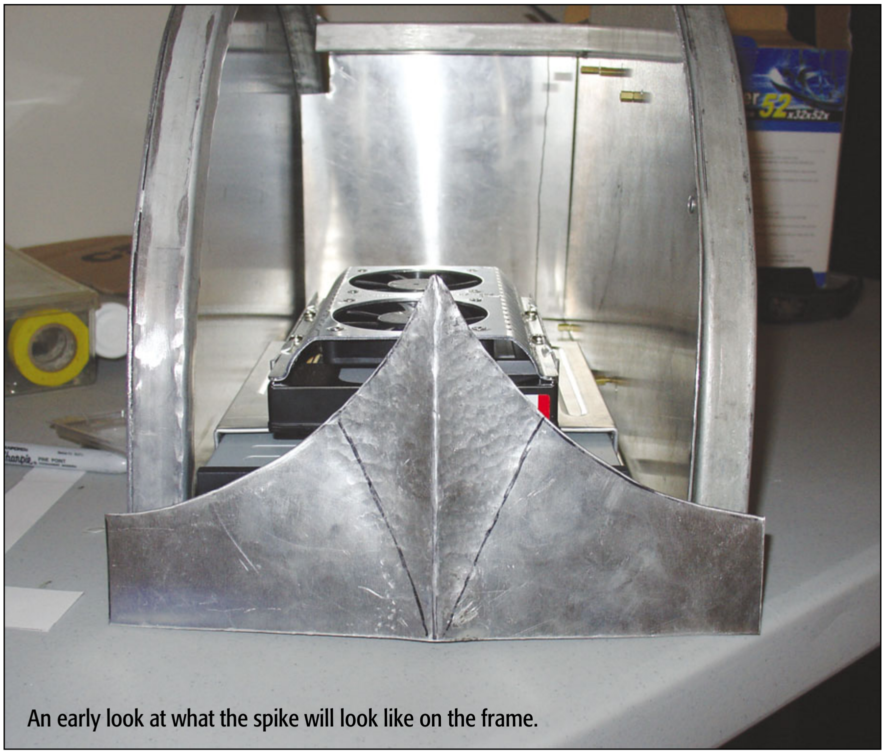

An early look at what the spike will look like on the frame.

where the backplane will be secured. This plate is the mounting for the power supply and the video card, and it also serves as a screen so that you can't see all the wiring in the back of the case when it's turned on.

With the motherboard pushed back an inch, it was no longer level with the cross member, so I had to come up with a creative solution that enabled me to still have the backplane in place and have it tie in to the cross member. After a little thought, I came up with a solution, which failed horribly. Originally, I decided to have the backplane curve along the arc of the frame and tie into the top of the frame. This way the backplane would be attached to the base and be curved at the top. In retrospect, this wasn't a very good idea, but I still forged ahead with the concept, thinking it would be a quick way to get out of a jam.

There's side-panel cuttin' to do here. Mounting the panel comes next.

Using a piece of ordinary printer paper, I made a template for the part I was going to cut. Everything was going smoothly at this point. I fit the part and got an idea of the radius. I then had to bend in the piece. After three hours of beating the metal in shape, I ended up with a mangled piece of metal that didn't fit well and looked terrible. It was time to find a new solution.

If I had taken an extra few minutes to think about what I was doing when I made the template, I would have realized that it would have been easier to bend an L shape in the top of the backplane. After butchering the curve concept, I decided to take a trip down the street to the metal shop and bend some right angles into this sucker.

A couple of hours later I had a backplane that fit the way I wanted and didn't look like it had been backed over by a car. I realize that I make mistakes (some sillier than others), and I don't always get things right the first time. So, sometimes I need a little help to get things done. The backplane still needs to have a mounting cut out for the power supply and graphics card, as well as a port hole drilled, but I've decided to hold off on that so I can finalize the placement of all the components and shorten the power supply leads.

Crossing My Wires

The power supply I'm using for this project is designed for use in a 1U case, but the lead wires are long enough to fit any case. I need to cut them down so that I can get the cleanest look possible. There's a lot of wiring to cut down, so I have a good few hours of soldering ahead of me. I also still need to replace the 8-pin connector with a 4-pin connector.

Wire management is a necessary evil in modding. I'm terrible at soldering, and I despise wiring. I would rather get hit in the head repeatedly than spend more than 20 minutes soldering, but with the open airflow architecture involved in this design, I need to reduce and minimize all the wiring in the case. It took a good five hours for me to cut down the wiring on the power supply and solder them all back together. I took about a foot off each of the wires and completely removed some of them.

I haven't even addressed the prominent spikes on the case yet, and for good reason. I will be adding these parts last with the help of Chris Crouse, who has also helped with the frame. During a recent visit, we talked about the look of the spikes and how they will fit together.

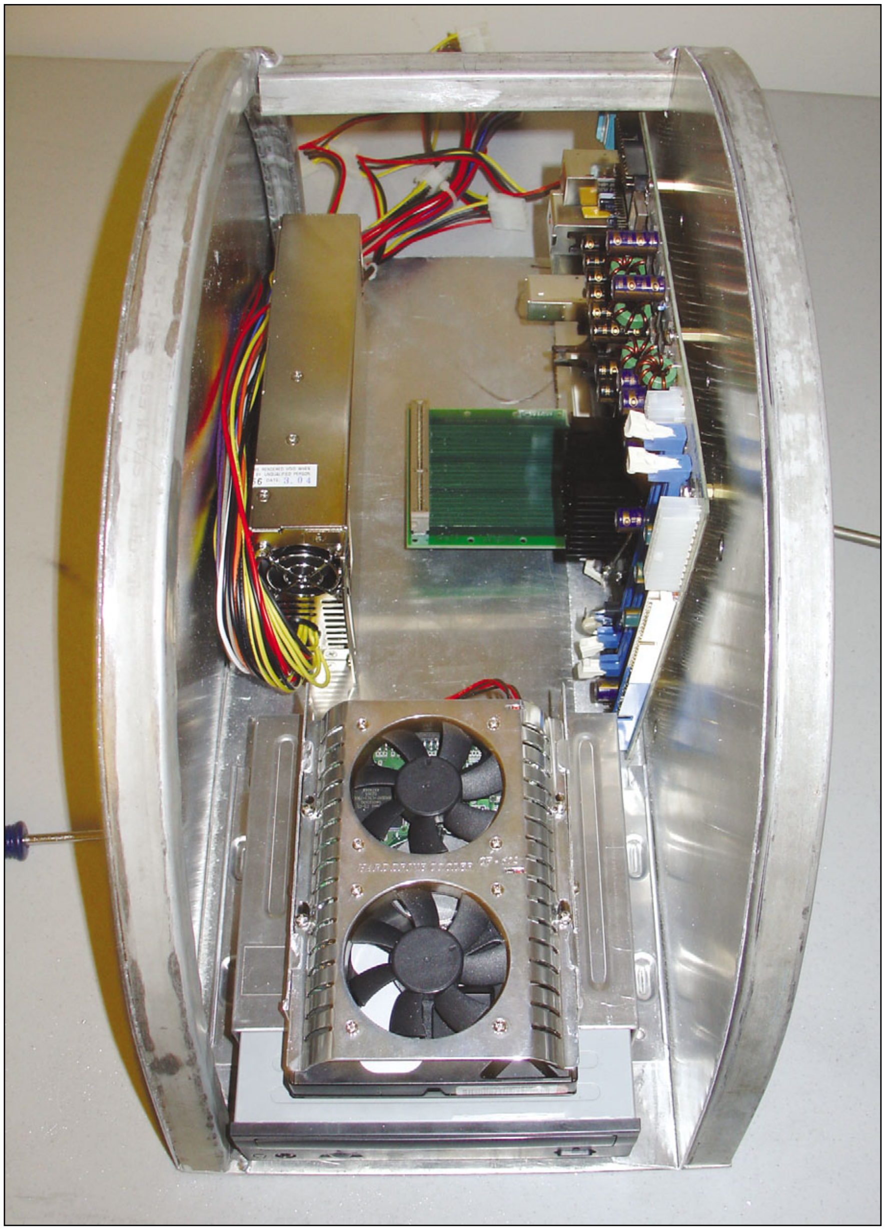

With the side panels mounted, Spike's insides are finally coming together.

After some experimentation with different gels and some innovation, Spike's lighting is coming along great.

Based on this, we decided to make a prototype spike.

Being that the bottom-front spike is really attached to the CD-ROM tray and not to the frame, I needed to get an idea of how I will mount the spike to the CD tray. By quickly mocking up the spike, I could figure out a mounting system that would be strong enough to hold the spike firmly in position but still allow the tray to move with ease.

We drew up a quick template, and Chris went to work forming the prominent ridge on the spike so we could see how the ridge looked on the spike. Even though we have the model to work from, I still want to experiment a little with the shape and form of the spikes. Originally, I had a ridge in the first drawing, but I removed it from the latest model. However, I like the medieval look that it gives the case, so I decided to add the ridges back in. The finished spikes will still be round and smooth, true to the model, but I won't be surprised if I change this look up a little.

On My Way

The frame of the case is done, and I'm well on my way to finishing this piece. However, I'm still far from done. The side panels, switch mountings, and the infamous spikes are the only major elements left to do. It has been a fun and frustrating time, considering the delays in getting material and the mistakes I've made. However, I've learned a lot from this project. I have also gained a new respect for aluminum; it's a lot harder to form and bend into shape than I originally thought.

The project has also been an altogether humbling one. I've been lucky to have the help of friends and craftsmen during this build, and I am eternally grateful for their time and patience. Although I'm far from done with this project, I'm hopefully through the worst of it and can enjoy smooth sailing to the finish. Then again, there have been some rough seas so far. I'm in open water, not knowing what lies in wait for me beyond the horizon. CPU

by Gareth Powell

Mad Reader Mod

This Will Get Your Motor Running

Give Us Your Mod

Have a computer mod that will bring tears to our eyes? Email photos and a description to madreadermod@cpumag.com. If we include your system in our "Mad Reader Mod" section, we'll send you a $1,500 Newegg.com gift certificate and a one-year subscription to *CPU*.

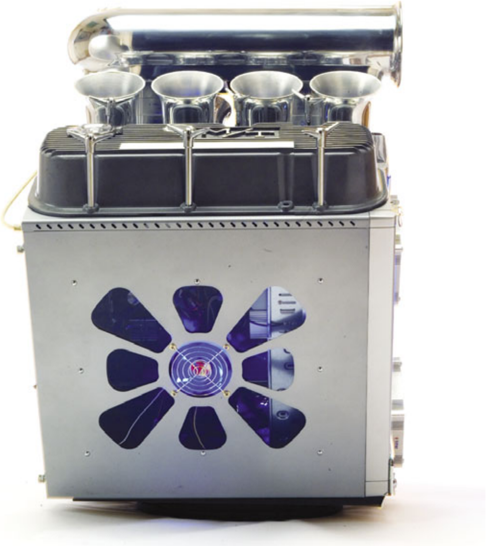

This is a case where it's what's on the outside that counts.

Vroom. With its scoop, Mickey Thompson valve covers, and velocity stacks, the V8 Engine mod is the master of any highway, "information" or otherwise. Dave Ledogar, aka Hotrodled, cooked up a plan to combine two passions, computers and cars, while restoring a 1967 Mustang. "I've been around cars since day one," Ledogar says. "I figured, let me put two cases together and see what happens."

Ledogar sketched plans for the mod, bought two identical Enermax cases, and then placed an old air cleaner, complete with velocity stacks, between them. To maintain a consistent appearance, Ledogar removed one of the case's fan-bearing panels, flipped it, and installed it on the

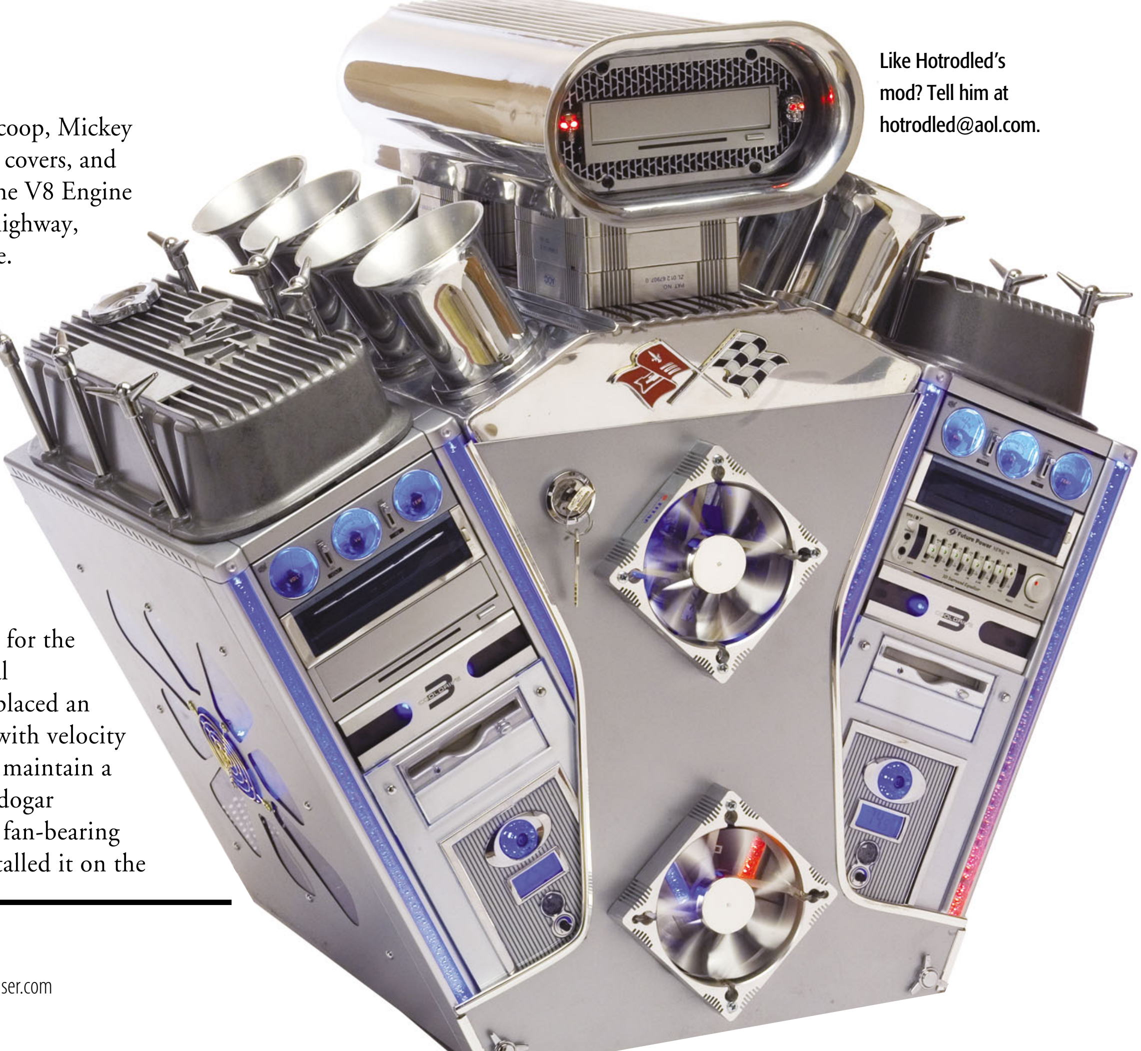

Like Hotrodled's mod? Tell him at hotrodled@aol.com.

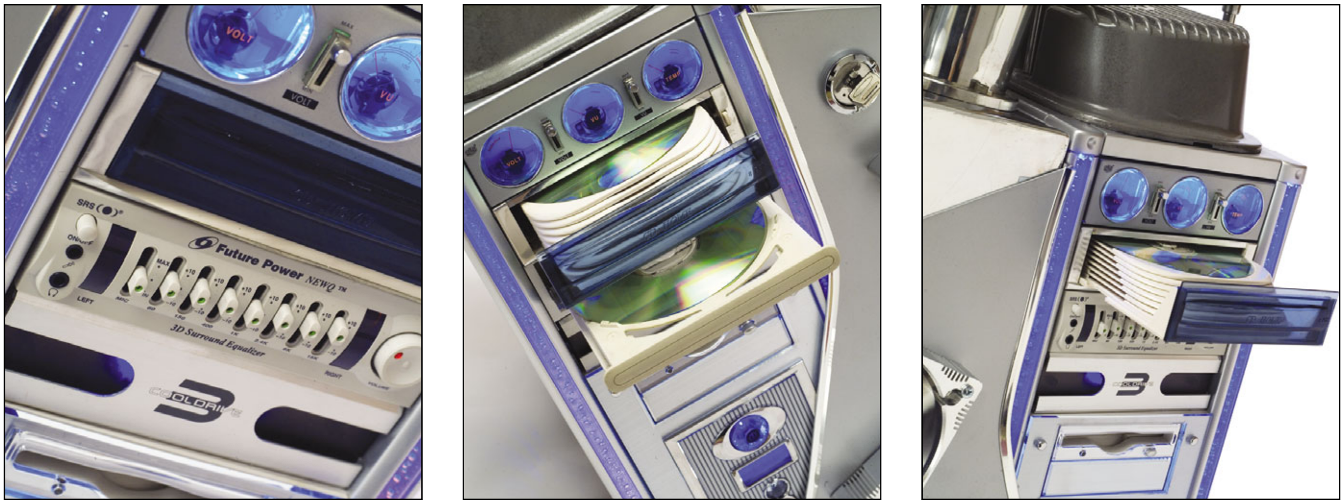

Both cases include 10 CD storage trays. You can retrieve the CDs by pushing the tray's button.

other side of the case. As a result, both sides of the engine sport the same, prebuilt window and case fan design. The front and back panels, which hold the cases together and support the engine's air cleaner, are the cases' former side panels. "You can't really cut it with sheers," Ledogar says of the side panel. "You have to use the Dremel tool with the cutting bit on it."

Next, Ledogar grabbed a friend's old valve covers, sandblasted them, and then placed them on the cases. Unlike the air cleaner, the long valve covers didn't quite fit on the tops of the cases, so he cut 3 inches from each cover and welded the pieces back together. Next, Ledogar drilled holes in the top of each case, attached the valve covers with custom chrome hold-downs, and added an oil cover.

The scoop, which houses the system's DVD-ROM drive, sits atop six 120mm Titan aluminum fans. Ledogar removed the fans from three of the cases to run cables from the DVD drive into the case. The same high-powered fans, which boast max airflows up to 80cfm each, suck cool air through the front panel and into the massive case. All told, the V8 mod has eight fans.

The ignition switch, also from a 1967 Mustang, presented Ledogar his biggest challenge. "It's a pretty heavy

Both Enermax cases have single-node temperature sensors, which monitor chassis and CPU temps.

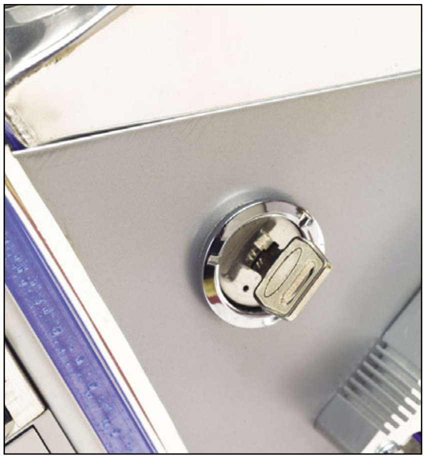

The V8 includes a real 1967 Mustang ignition, so it offers the same resistance as a car's ignition. When you turn the key, the computer emits an engine sound and then powers on.

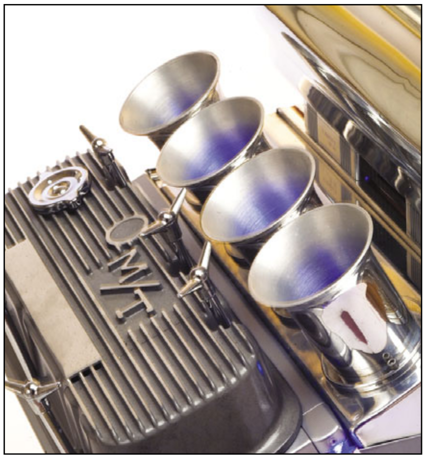

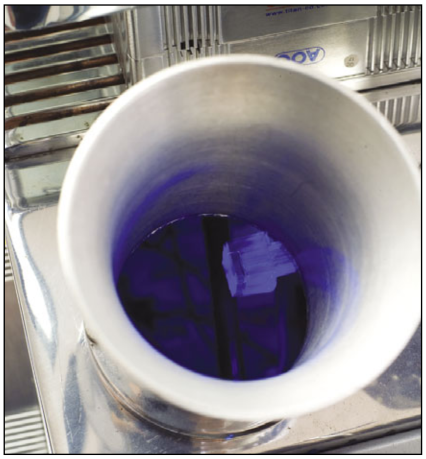

A long cold-cathode light gives the velocity stacks an eerie glow.

unit to sit on the front door, so I had to make some supports," Ledogar says. He built a cage to support the switch and then wired the switch to the system. Of course, the engine wouldn't be a V8 if it didn't sound like one. When you turn the system's ignition switch, the computer's surround sound speakers roar to life. A second turn powers on the computer.

A 550-watt Antec power supply supports the power-sucking beast. Inside, the system sports a 1.6GHz Intel P4, 512MB of RAM, two Maxtor hard drives, and a GeForce FX 5200 graphics card. The guts don't live up to the chrome-covered chassis, but Ledogar won't have any trouble upgrading the components, thanks to the $1,500 Newegg.com gift certificate we sent him.

We figure that Ledogar can pump up the system with a 939-pin 2.2GHz AMD Athlon 64 3500+, ASUS A8V mobo; 256MB Gigabyte ATI Radeon X800, and a 1GB Corsair XMS PC3200 DDR RAM twin pack. That ought to last a few miles. ▲

---

by Joshua Gulick

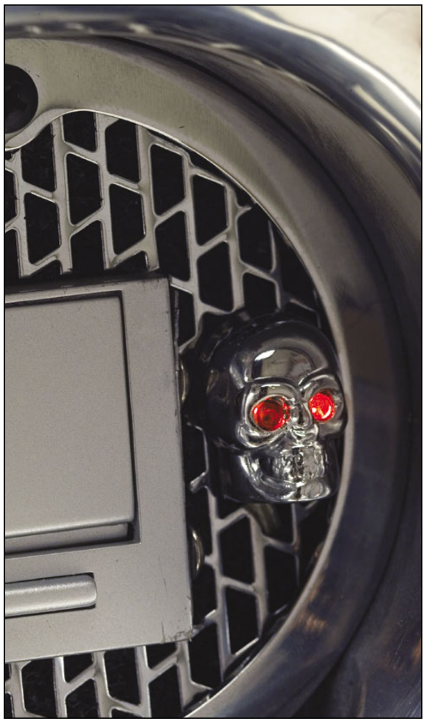

Thanks to the PSU's 12V cables, even the extras draw juice.

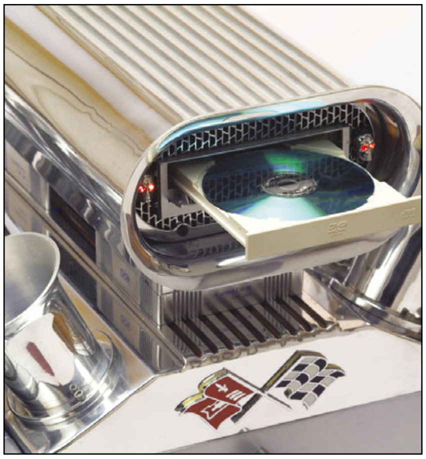

Yep, the DVD-ROM is functional. The drive's cables drop through three empty case fan shells into the computer. Ledogar spotted the flag emblem on eBay.

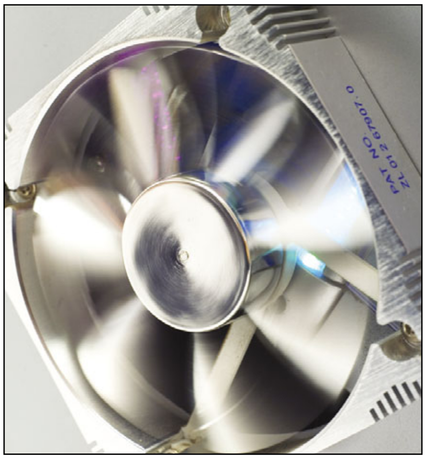

Dual 120mm front-panel fans force-feed cool air into the bowels of Ledogar's beast.

3com

The Company That Wired the Network Now Unwires it

Introducing the 3Com® Suite of 802.11 Wireless Solutions

Imagine untethered network and Internet access from anywhere in your home or office. 3Com® brings freedom of choice to wireless with one of the most comprehensive wireless networking suites on the market today.

3Com wireless products come in a variety of price points and are scalable from home and small office to enterprise solutions. Our OfficeConnect® line for home and small offices is among the first to ship with the latest 802.11g standards and new 256-bit Wi-Fi Protected Access (WPA) encryption security.

Our enterprise-class products include an 11a/b/g PC Card with patented XJACK® antenna and dual-mode upgradeable access points that offer you the freedom to mix and match 802.11 a, b, or g standards to meet different coverage and bandwidth needs, and featuring the latest suite of security standards enabling customers to match the level of security to their specific environments.

Best of all, every 3Com wireless product is backed by 30 years of Ethernet networking expertise.

Feel free to unwire your office. Visit www.3com.com/shop today to find the 3Com wireless solution that works best for you.

3Com® OfficeConnect® Wireless 11 a/b/g PC Card, 11g Gateway

3Com® 11a/b/g Wireless PC Card with XJACK® Antenna

3Com® Wireless Access Point 8750

Tri-mode. Upgradeable.

To find a reseller nearest you, go to: www.3com.com/locate_reseller

CDW Insight PC Connection PCMall

Copyright © 2003 3Com Corporation. All rights reserved. 3Com, the 3Com logo, OfficeConnect and XJACK are registered trademarks of 3Com Corporation. All other product names may be trademarks or registered trademarks of their respective companies.

Advanced Q&A Corner

Get informed answers to your advanced technical questions from *CPU*. Send your questions along with a phone and/or fax number, so we can call you if necessary, to q&a@cpumag.com. Please include all pertinent system information.

Each month we dig deep into the mailbag here at CPU in an effort to answer your most pressing technical questions. Want some advice on your next purchase or upgrade? Have a ghost in your machine? Are BSODs making your life miserable? CPU's "Advanced Q&A Corner" is here for you.

BioMed asked: I'm building a custom PC and have a question regarding RAM. I'm thinking about buying an ASUS P4P800 motherboard with a Pentium 4 that has a 533MHz FSB. I also want to purchase some PC3200 DDR400 SDRAM but am not sure this is a good idea. If the RAM is running at only 400MHz and not 533MHz like the CPU, will this cause a bottleneck and increase the time the CPU is waiting for information?

A: It's very easy to be confused by "effective" and "actual" clock speeds in a situation such as this. A Pentium 4 with a 533MHz frontside bus is actually running with its FSB clocked at 133MHz, but because the NetBurst microarchitecture of the CPU is "quad-pumped," the effective speed is 533MHz (133.3MHz x 4). The same holds true for the memory. DDR stands for double data rate, which means that two bits of data can be transferred per clock cycle. Due to this fact, memory that is actually clocked at 200MHz has an effective speed of 400MHz (200MHz x 2), hence the DDR400 moniker. The PC3200 label designates the memory's maximum theoretical bandwidth of a single channel of DDR400 RAM (400MHz x 8 = 3.2GBps).

The Pentium 4 rides along on a 64-bit bus, so using the same equation (effective clock speed x bus width in bytes), we find that when clocked at 533MHz, its peak bus bandwidth is just a bit below 4.3GBps. In this case, a single channel of 400MHz of PC3200 RAM would not be able to saturate the P4's bus, causing a bottleneck. This is where the i865 chipset powering the ASUS P4P800 motherboard comes in, however. The i865 chipset has a dual-channel memory controller, which means each channel can pump a theoretical maximum of 3.2GBps of data per clock cycle for a grand total of 6.4GBps. If you configure the system to work in dual-channel mode with PC3200 RAM, it'll be perfectly capable of feeding adequate amounts of data to the CPU to keep the pipeline filled. Just keep in mind that you'll have to buy two matched sticks of RAM to take advantage of the chipset's dual-channel memory controller when you make your purchase.

Kin Ho asked: I recently bought a barebones SFF box. The setup went perfectly until I tried to install WinXP Pro to boot with a SATA hard drive. The following are the specs of my system:

Barebones model: Biostar iDEQ 200A
Motherboard: ATI 9100 IGP
SATA RAID chipset: VIA VT6420, with setting utility v2.21
CPU: Pentium 4E Prescott, 2.8GHz, 1MB L2 cache, 800MHz system bus
RAM: Mushkin 184-pin SDRAM DDR400 (PC3200) 512MB x 2
Optical disc: Lite-On DVD-ROM/CD-RW combo
AGP add-on: None
Hard drive tried: Maxtor 80GB SATA

WinXP installation halted and gave me a screen saying, "Can't find any master storage device. Installation can't continue." I tried using F6 and a diskette with the correct VIA VT6420 SATA RAID driver during the initial WinXP installation but got the same error message. I then emailed Biostar techs, and they gave me another copy of the SATA driver, but I got the same error. They had no further suggestions. Is this a sign that this technology is not yet mature in the computer industry?

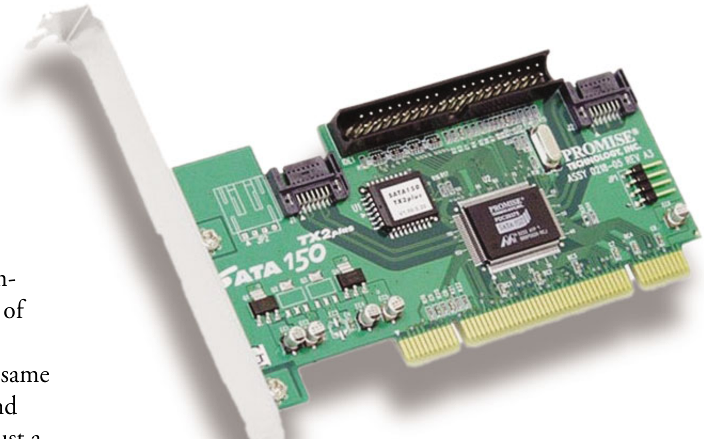

You can use cards such as the Promise FastTrak S150 TX2plus RAID 5 controller to add SATA and RAID capabilities to a system that lacks these features.

A: Without question, Kin, Serial ATA controllers and SATA hard drive technology is still a bit rough around the edges in terms of compatibility and maturity. While your problem most likely stems from the fact that the SATA controller in your system is not natively supported in WinXP and it needs a driver before the OS can even see the target install drive, we've also seen many cases where certain SATA controllers/motherboards just don't play well with certain SATA hard drives. Sometimes this is a BIOS-related issue; other times it can be drive firmware-related.

More than likely, though, you are experiencing either a conflict between devices on your SATA channels and devices on your standard ATA channels or the driver files you are trying to feed WinXP at the initial install screen aren't being recognized. If you've made the driver diskette by extracting files downloaded from the manufacturer's site, make sure there is a file called Txtsetup.oem on the diskette. This file is an identification and configuration file that tells your computer what driver it is installing and what files to call during installation. When you inserted the driver diskette and pressed ENTER, did a box come up on your screen that gave the name of the driver you were installing? If you didn't see this, that means the system couldn't find the driver on your diskette, and that's most likely because all the files, especially the ID file, weren't in the right place.

The other potential issue is that your CD-RW drive (which you might want to jumper as a slave device regardless, just to rule out a master/slave configuration conflict), could be conflicting with the primary or secondary SATA channel on your motherboard. Your BIOS menu, depending on the manufacturer, could have a setting called "enhanced mode," which should allow you to run any combination of SATA or PATA (legacy Parallel ATA) drives on any channel. However, if you want to rule this issue out completely, try unplugging your CD-RW drive from its IDE channel altogether and see if that allows your main hard drive to be recognized. By the way, Maxtor SATA drives have a pretty good track record for compatibility with various SATA controllers, so hopefully that isn't your issue. Good luck!

Bill T asked: My computer has two SPDIF output ports. What does SPDIF stand for? What are these ports used for? What kind of hardware device can I hook up to these ports?

A: The acronym SPDIF stands for Sony/Philips Digital Interface, which is a standard audio file transfer format. It specifies the transfer of digital audio signals from one device to another without having to be converted first to an analog format. With a pure digital connection, there is no need to convert from digital format to analog, a process that introduces some level of signal degradation and loss.

You can make SPDIF connections to a pair of digital-ready speakers, such as the ever-fabulous Logitech Z-680s, with either an SPDIF coaxial or fiber optic cable. This type of connection will allow for the purest reproduction of a digital audio source on your computer, such as a CD, DVD, and MP3.

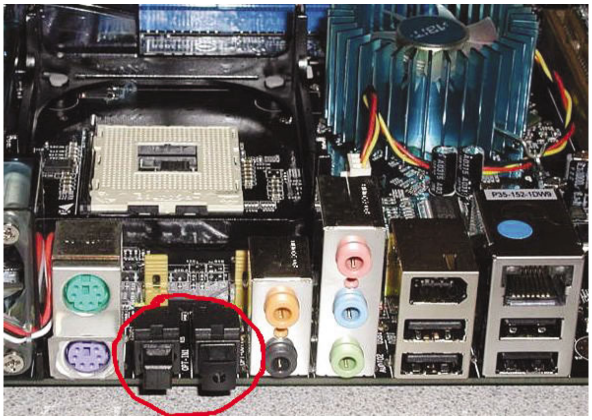

SPDIF (Sony/Philips Digital Interface) allows for a pure digital connection to a set of digital-ready speakers, from your sound card or onboard sound, with a fiber-optic or coaxial cable.

Vincent Ngo asked: I just bought a new Promise FastTrak S150 TX2plus RAID controller with two 120GB Western Digital SATA Special Edition hard drives. I've set them up in a RAID 0 (striped) configuration and use the system mainly for video editing. What stripe and NTFS cluster size should I use to format the array to get the best performance?

A: When formatting the array, the cluster size should match the stripe size, but the "best" stripe and cluster size to use with a two-drive RAID 0 configuration has been the topic of much debate. You'll likely get a different answer to this question depending on whom you ask and how they use their system on a daily basis. Enthusiasts who spend most of their time gaming, browsing, and benchmarking will likely say smaller stripe/cluster sizes offer the best performance, but ask someone who constantly works with huge files (greater than 300MB) all day, and he'll probably say the larger the stripe/cluster size, the better. The fact of the matter is that both answers are correct, but believe it or not, the performance differences are quite small.

We have been experimenting with a pair of 74GB Western Digital Raptors set up in a RAID 0 configuration for quite some time and found that varying the stripe size doesn't impact real-world performance all that much. Benchmarks such as HD Tach and SiSoft SANDRA's File System Performance module did show slight performance deltas. But when we timed how long it took to copy large files (three 700MB ISOs) to the array, the entire process was only a few seconds faster when we used a larger stripe size (16KB stripe vs. 64KB stripe). The fact that you're editing video leads us to believe you're working with big files, so a larger stripe size may yield a slight performance gain, but don't bother reformatting if you've already configured the array with a smaller stripe size. You're probably not going to notice the difference.

What you may want to consider is that a larger stripe/cluster size is a less efficient use of drive space. When using a 64KB stripe/cluster size, files that are smaller than 64KB will still use up 64KB of hard drive space. If using your hard drive space more efficiently is more important to you than saving a few seconds here and there, formatting the array with a smaller stripe/cluster size will be the better choice. ▲

by Dave Altavilla and Marco Chiappetta,
the experts over at HotHardware.com

For more answers to your questions, subscribers can see www.cpumag.com/cpuaug04/q&a.

BIC-TCP

Giving Researchers High-Speed Hopes

The history of TCP (Transmission Control Protocol) shows it's an extremely useful protocol for transmitting data, even over lines with heavy traffic. But long-distance transmissions cause trouble for TCP in the form of long delays over high-speed networks.

With several high-speed, large-bandwidth networks now employed around the world, researchers have worked in recent years to find a new protocol for making full use of the available bandwidth. For example, networks such as ESNet (Energy Sciences Network) offer bandwidth up to 10Gbps when connecting research facilities around the world.

Because these high-speed networks often connect remote locations worldwide, the problems with RTT (round trip time; the time required for a packet to make a trip along with the return response) delays are magnified by the distances the data must travel. Although researchers have introduced and proposed several new protocols in recent years (see the "High-Speed Protocols Compared" sidebar), one of the most promising appears to be BIC-TCP (Binary Increase Congestion Control-TCP). BIC-TCP received the highest marks in a recent series of tests from Stanford researchers on high-speed protocols.

BIC-TCP's Proposal

The trouble with TCP is that it has difficulty filling the bandwidth that's available in a high-speed network (a problem that's usually called congestion control). TCP does increase its congestion window steadily each time it has no packet loss, but the increases are minimal enough that TCP has trouble reaching the full available bandwidth in a desirable time. If it begins encountering errors and packet

Binary Search Increase

The researchers who developed BIC-TCP (Binary Increase Control-Transmission Control Protocol) incorporated binary search increase in the protocol. Binary search increase provides reliable feedback on any network congestion and lost packets, allowing BIC-TCP to aggressively increase its transmission speed toward the maximum allowed by the high-speed network. Binary search increase defines the **minimum window size** as the speed at which no packet loss occurs. The **maximum window size** is the top transmission speed allowed. The **target window size** is the midpoint between the minimum and maximum. (The **current window size** is the actual transmission speed and attempts to climb to the target window size.) Binary search increase calculations cause all of the window sizes to continually change.

1 As the transmission begins, binary search increase calculations allow the current window size to aggressively increase toward the target window size because the difference between the minimum and maximum is large.

2 Once the current window size reaches the target without any packet loss, the minimum window size increases to match the previous target and the target window size increases to the midpoint between the new minimum and maximum (which remains steady). The current window size then continues increasing toward the new target.

3 When a packet loss occurs, the maximum window size drops to match the position of the current window size where the packet loss occurred. The minimum window size drops to match the location where the current window size drops after the packet loss. The target window size is readjusted to the midpoint between the new minimum and maximum. Then the current window size restarts its climb.

4 As the current window size continues to increase without packet losses, binary search increase continues its calculations for the minimum and target window sizes. As the difference between the minimum and maximum window sizes shrinks, binary search increase forces the rate of increase in the current window size to become less aggressive, which allows for fewer packet losses, despite the increase in transmission speed.

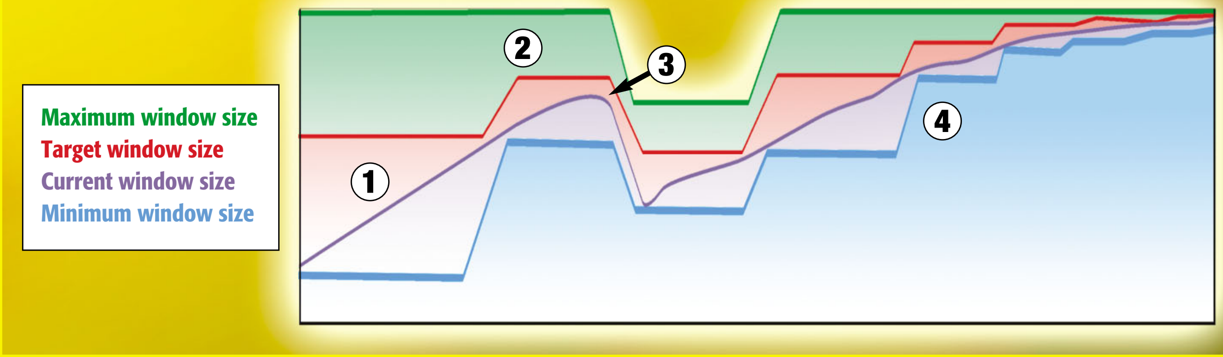

Graphics & Design: Carrie Benes & Jason Codr

TCP's Weakness

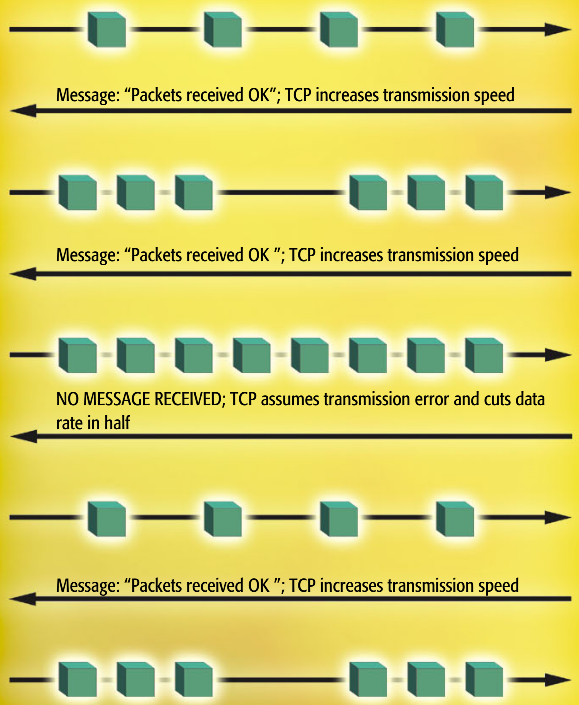
Message: "Packets received OK"; TCP increases transmission speed
Message: "Packets received OK "; TCP increases transmission speed
NO MESSAGE RECEIVED; TCP assumes transmission error and cuts data rate in half
Message: "Packets received OK "; TCP increases transmission speed

TCP is a key component in transmitting data across computer networks, especially the Internet. Some estimates say TCP accounts for 90% of all data transmission across networks. TCP is extremely reliable because the sending medium always receives notification of the successful transmission of the packets. The reliability factor works especially well on the Internet, where congestion and lost packets are common.

When TCP receives successful notification of the receipt of packets, it slowly increases the transmission speed. However, if TCP doesn't receive a receipt message, it assumes lost packets and network congestion and automatically halves its transmission speed. It then restarts the process of slowly increasing transmission speeds whenever receiving notification of a successful packet receipt.

Most Internet users never notice when TCP slows down to deal with congestion problems because the protocol still works much faster than the typical Internet user's connection speed. TCP will be a reliable, successful protocol for most Internet users for the foreseeable future.

However, in a high-speed network (usually more than 1Gbps in bandwidth), TCP's delays are debilitating. TCP simply cannot take full advantage of the high-speed network's available bandwidth. TCP increases its transmission speeds far too slowly to ever take full advantage of the available bandwidth. Researchers have developed several different Advanced TCP options (such as BIC-TCP) that are more aggressive about increasing transmission speeds while remaining able to work alongside TCP.

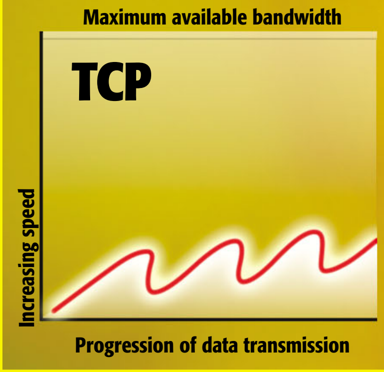

Because of its emphasis on providing reliability, TCP cannot meet the high-speed transmission needs of a large network. As shown at left, TCP's transmission speeds cannot climb quickly enough to meet the maximum because the protocol drops the speed by half each time it encounters an error. Advanced TCP options (shown at right) are more aggressive about increasing transmission speeds, backing off when encountering errors but quickly resuming an aggressive increase in transmission speed.

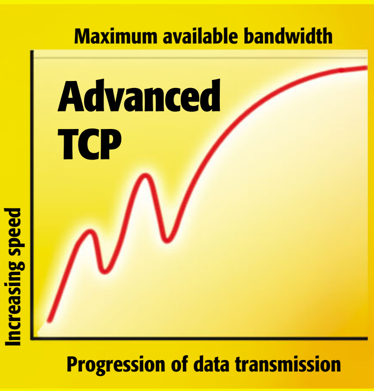

losses, TCP's bandwidth problems become magnified.

Researchers can fix such bandwidth problems with newer protocols designed specifically for high-speed networks. However, one of the biggest concerns researchers have had about some of the high-speed network protocols involves how they work alongside TCP (usually called TCP friendliness). As it works with TCP, the new protocol must treat TCP fairly in the amount of bandwidth it occupies. In other words, while scaling to utilize the full bandwidth of the high-speed network, the new protocol must avoid overwhelming TCP's share of the bandwidth.

The researchers who developed BIC-TCP at North Carolina State University's Department Of Computer Science (Injong Rhee, Lisong Xu, and Khaled Harfoush) had another serious concern about the other available protocols: RTT fairness. When different data flows with varying RTTs occur within a high-speed protocol, some data flows consume more than their share of bandwidth, causing problems. The North Carolina State researchers are hoping they've addressed all of these concerns with BIC-TCP.

When creating BIC-TCP, the researchers put together two important components: Binary search increase and additive increase.

Binary Search Increase. Binary search increase affects the method by which a protocol works to achieve its target sending rate, also called the window. (See the "Binary Search Increase" graphic.) As the protocol works to find the window size that maximizes its bandwidth, the protocol constantly adjusts the current, target, maximum, and minimum window sizes to meet the feedback it receives concerning packet loss.

With TCP, this move toward the maximum available bandwidth occurs far too slowly to achieve the desired results on a high-speed network. When using BIC-TCP, binary search increase allows the protocol to be aggressive when moving from the current window size to the target window size whenever the difference between the two is large. As the difference between the two window sizes shrinks, binary search increase forces the protocol to be less aggressive, which has the effect of treating TCP fairly.

Additive Increase. Additive increase is a process by which TCP sends additional packets per RTT, as long as the previous set of packets arrived without error. BIC-TCP's use of additive increase makes it more stable.

High-Speed Protocols Compared

In addition to BIC-TCP (Binary Increase Control-Transmission Control Protocol), several researchers have developed other high-speed protocols–often called Advanced TCP–in recent years. All of these protocols improve on TCP by increasing their congestion windows more quickly than TCP can increase them. These protocols base the rate by which they increase the congestion window size on the overall size of the congestion window, which allows the protocols to maximize the data rate far more quickly than TCP.

Fast TCP. Developed at the California Institute Of Technology, Fast TCP measures a different factor (queuing delay instead of loss probability) to determine congestion in the network. By measuring this factor, Fast TCP can increase its congestion window more quickly than can TCP. Fast TCP was involved in a key test of Advanced TCP over a transatlantic network in 2002.

Hamilton TCP (H-TCP). Developed at the Hamilton Institute in Ireland, H-TCP works at balancing the increased throughput speeds of the high-speed network while maintaining the stability of TCP. H-TCP is aggressive about increasing the congestion window after it reaches the target threshold, therefore more quickly responding to changes in available bandwidth.

High Speed TCP (HS-TCP). HS-TCP uses a table to determine by what factor to increase the congestion window. HS-TCP is similar to H-TCP in its aggressiveness about increasing the congestion window when the available bandwidth increases. However, H-TCP uses an algorithm to determine the rate of increase instead of a table.

High Speed TCP Low Priority (HSTCP-LP). Developed at Rice University, HSTCP-LP only focuses on using the excess network bandwidth instead of attempting to determine its fair share of the bandwidth. Developers of HSTCP-LP borrowed some aspects of HS-TCP.

Scalable TCP (S-TCP). Developed at the Communication Engineering Lab of Cambridge University in England, S-TCP slightly tweaks TCP, allowing it to increase its congestion window exponentially rather than using a linear increase. Because the changes to TCP are minimal, S-TCP researchers say their protocol works extremely well alongside TCP.

The SLAC (Stanford Linear Accelerator Center) recently published its findings on tests it performed on many Advanced TCP options. The SLAC tested each protocol for reliability, fairness, and performance over long-distance networks. SLAC researchers called BIC-TCP the best Advanced TCP performer in its tests, scoring well across all categories. Other protocols, including S-TCP and Fast TCP, also received high marks but had a few minor problems. Researchers say S-TCP was too aggressive in taking bandwidth from TCP and Fast TCP experienced some problems with slowdowns during periods of high reverse traffic. ▲

There's No Comparison

One confusing aspect of BIC-TCP has been the widely reported comparisons to DSL that accompanied the recent announcement of the publication of a BIC-TCP research paper. However, because DSL and BIC-TCP are two different items—DSL is a medium and BIC-TCP is a protocol—such comparisons are inaccurate. A protocol actually carries the data across the medium.

Although it's impossible to compare the speed of BIC-TCP with DSL, researchers say BIC-TCP will allow a 10Gbps network to use its full capacity, something TCP alone cannot do. Researchers say that BIC-TCP will provide more benefit to any high-speed network beyond 1Gbps than TCP can. Researchers estimate that over a 10Gbps, long-distance network, TCP only will be able to reach speeds of 450Mbps. When running alongside TCP, BIC-TCP would occupy the remaining bandwidth of the network.

Another source of confusion concerning BIC-TCP concerns how much it will help the average user right now. The answer is not much. Because the vast majority of users connect to the Internet through a cable connection, a DSL connection, or even a dial-up connection, TCP can easily deliver data at the maximum possible bandwidth. TCP's weaknesses really don't show up in a significant way until the protocol is applied to a high-speed network, usually at least 1Gbps. Therefore, most ISPs won't be employing BIC-TCP (or any other high-speed protocol, for that matter) for use with current consumer Internet connections.

Scientists and researchers, however, are ready for the high-speed protocols, which will lead to transmission of petabytes (1 million GB) and exabytes of data (1 billion GB). High-speed protocols will open avenues of research previously thought impossible because of data-transfer limitations.

by Kyle Schurman

IT'S SHOWTIME

DIGITALMOVIE™ DVD

Talk about box office hits. Verbatim's new DigitalMovie™ DVD recordable media is so hot you can almost smell the popcorn. Available as DVD-R or DVD+R, this 4X speed DVD media looks – even feels – like a Hollywood movie reel. But don't think you need a director's chair to use them. These mini-reels are perfect for home movies, business presentations, transferring files – anytime you want to add a little fun to the show. And of course, DigitalMovie DVD has the same unsurpassed performance and reliability you've come to expect from Verbatim. So make your next show a feature presentation. DigitalMovie DVD from Verbatim. It's the reel thing. For more information, contact us at 800-421-4188 or online at www.verbatim.com.

Verbatim.

PUT YOUR WORLD
ON VERBATIM

© Verbatim Corporation 2004

Alienware ALX X2

Double Your Pleasure, Double Your Fun, Double Your Video Cards

Few technologies receive unanimous, wide-ranging accolades when they're introduced. No matter how good a technology appears initially, there are always a few skeptics around to throw cold water on the idea. But when Alienware announced its new ALX (Alienware Luxury Experience) line of high-end-graphics desktop PCs at the recent E3 conference, the skeptics were as scarce as . . . well, sightseers inside the fence at Area 51.

Alienware's ALX computers will feature several impressive technologies, including dual video cards, dual CPUs, and factory overclocking. Alienware's new Video Array technology and proprietary X2 motherboard will make everything work together. The company expects its ALX line to allow gamers and high-end-graphics artists to run all games and applications to the highest resolutions and power settings with no loss in performance.

Twice As Nice

The ALX will be a unique machine because Alienware will not sell its technologies to third-party manufacturers. Neither will the company sell the X2 motherboard separately from the ALX computer, meaning enthusiasts who want to build their own computers won't be able to use the X2 to construct their own dual-processor, dual-video card systems.

The ALX can use almost any video cards from leading manufacturers, as long as those cards are the same model. During its E3 demonstration, for example, Alienware used dual NVIDIA GeForce 6800 Ultra graphics cards. The ALX will use Xeon processors from Intel, but you don't have to get a double-CPU ALX; Alienware will also offer a single-CPU version.

The first Aurora ALX computer models began selling at www.alienware.com/alx in early June at a base price of $4,799. But the initial ALX Aurora models have only one video card and one CPU, and for now they won't feature PCI-E (PCI Express) or Alienware's Video Array technology.

The dual-CPU, dual-video card models (which will probably be called ALX X2) with PCI-E and Video Array should appear in the fourth quarter of 2004. Alienware won't release the high-powered ALX X2 models until Intel's Turnwater chipset becomes available. Alienware officials also expect the X2 motherboard to be the first to support dual PCI-E graphics slots. Pricing for the high-end ALX X2 computer, which will only be sold directly through Alienware, hadn't yet been set when we went to press.

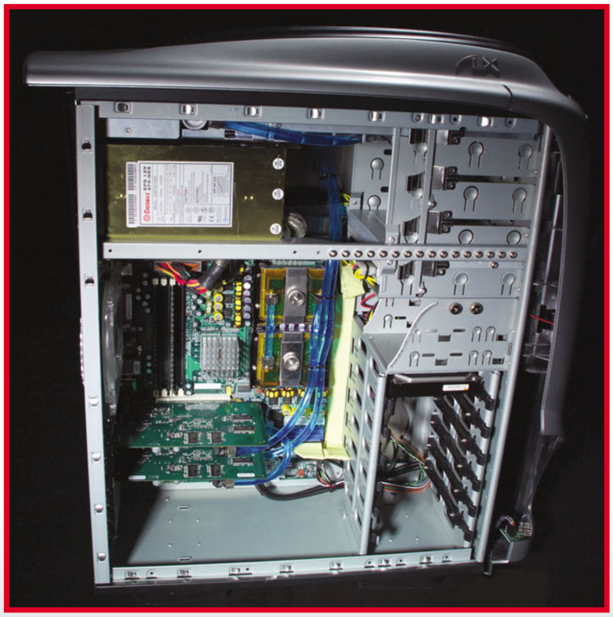

In this photo of the interior of the ALX X2, you can see the dual video cards in the lower-right corner of the case, the dual processors in the middle, and the liquid-cooling system with the blue tubes running throughout the case.

Video Array Technology

Alienware's Video Array technology makes the dual-video card setup work. Video Array adds a proprietary Alienware software layer and a piece of hardware, called a Video Merger Hub, to the basic graphics rendering process. These additions allow the dual video cards to maximize their graphics rendering power and work together smoothly. (See the "Video Array Mechanics" graphic.)

Alienware has developed Video Array to be as simple to use as possible. Video Array can use off-the-shelf video cards and almost any kind of graphics chipset. No special graphics drivers are required, either, which will be handy for gamers who often need to swap video card drivers.

By keeping the parameters for Video Array simple, Alienware will allow users to make changes to their ALX system video cards as graphics technologies improve and new video cards debut. Video Array supports both OpenGL and DirectX 9, as well. However, the company will customize all the other components in the ALX, such as the chassis, motherboard, power supply, and cooling system. Users will have a very difficult time swapping out these parts.

Old Ideas, New Technologies

Alienware's dual-video card setup isn't a completely new concept: 3dfx Interactive introduced a related idea, called Voodoo2 SLI (scan-line interleave), several years ago. However, Voodoo2 SLI's two video cards created the image by drawing alternating lines on the screen. (Video Array has one video card draw the top half of the screen and the other card draw the bottom half, with adjustments made to each card's workload as the scene demands.) SLI never really caught on, in part because newer video cards quickly surpassed what Voodoo2 SLI could do and in part because of Voodoo2 SLI's high cost. Other manufacturers have tried placing more than one graphics chip on the same video card with little success in the marketplace.

The idea behind it might be old and simple, but the Video Array technology definitely is new and complex. Alienware's Video Array technology isn't based on any intellectual property from other companies, according to company officials. The company worked on Video Array for nearly three years before introducing it at E3.

"Alienware was able to keep the news about this technology from leaking to the public as a result of having our own in-house development team," Alienware CTO and executive vice president Humberto Organvidez says. "This limited the sharing of information only to the core team that was actively working on the Video Array and X2 technology."

Alienware's ALX X2 Key Features

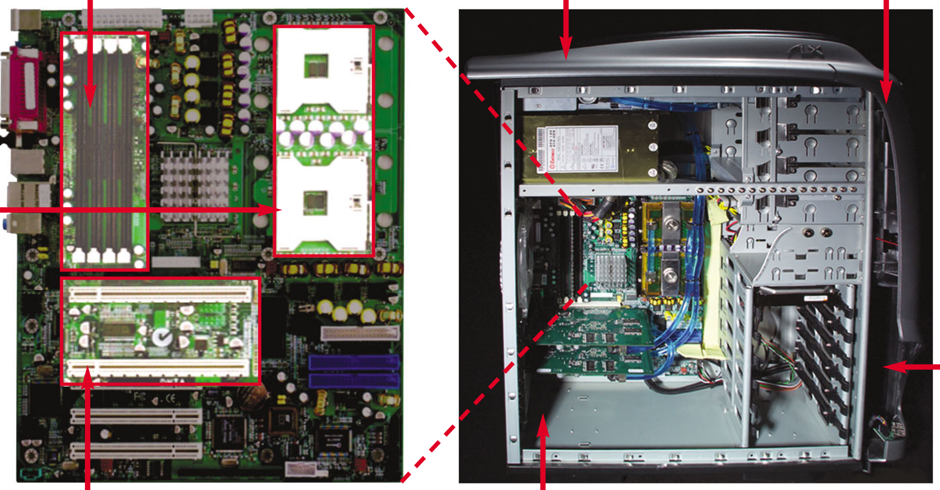

ALX X2 models will use up to 8GB of DDR2 memory, although the initial model, the ALX Aurora, is offering 1GB or 2GB of memory.

A large, customized power supply—at least 800W—will be located here in the case. It needs to be large to provide enough power to the dual video cards and the dual CPUs. Because the power supply is so large, it can provide power to the cooling system, as well, eliminating the need for a second power supply.

An LED located here shows the case temperature.

Those purchasing the ALX X2 will have the option of one or two CPUs. The dual-CPU setup should help prevent problems with CPU lag that could occur because of the high graphics-processing speed of the two video cards.

The ALX X2 designers reworked the size and configuration of the hard drive cages, located here, giving users easier access to the drives for future upgrades.

Alienware's X2 motherboard will be the first to feature two PCI-E x16 graphics slots. The company has a patent pending on its X2 motherboard design. The dual video cards must be the same brand and model, but no special video card drivers are required to make the Video Array technology work correctly. The two PCI-E slots open several interesting possibilities, such as using two dual-head video cards, which could combine to support four-monitor display output.

The ALX X2 case design will provide maximum cooling efficiency with minimal noise, thanks mainly to its copper-plated, liquid-cooling system that will provide cooling blocks for both the dual video cards and the dual processors. The liquid-cooling system was designed exclusively for the ALX X2 and will have components running throughout the case. The ALX X2 case will also have ultra-quiet, 120mm fans and plenty of fins to allow for a large amount of air circulation.

Video Array Mechanics

Alienware's Video Array architecture is vastly different from standard graphics processing architectures that use one video card. Alienware uses a software layer between the application in use and the video card drivers. The Alienware software layer determines how to most fairly divide the rendering workload. The video cards then draw the portion of the screen the Alienware software has assigned to them. Finally, the Alienware Video Merger Hub, which sits between the video cards and the output monitor, combines the two partial images. The Alienware software layer and the Video Merger Hub both have patents pending.

The rendering load is split 50/50 between each video card as Video Array begins working. As part of the technology, Video Array measures after each frame how long each video card needed to render its portion of the screen. Video Array then has the software layer change the division of the work, based on its measurements, to give the video card that rendered its portion of the previous frame more quickly more work to do in the next frame. This logic, called predictive load balancing, attempts to maximize the graphics processing power from each video card. For example, one frame may involve a simple sky with only a few polygons in the upper half of the screen with a complex scene in the bottom half of the screen, requiring thousands of polygons. In this example, predictive load balancing calculations could result in one video card drawing 85% of the screen and the other video card drawing 15%. The ratio of the workload constantly changes in real-time as the Video Array technology applies predictive load balancing.

Standard Graphics Architecture

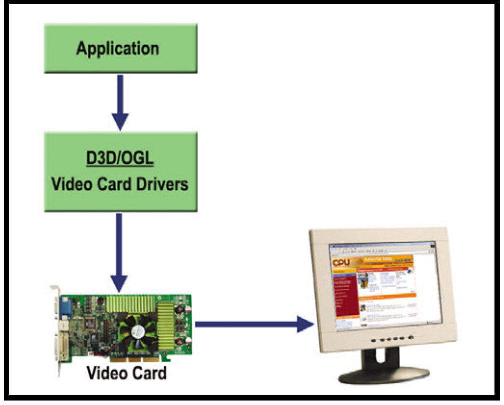

Alienware-Developed Software Layer

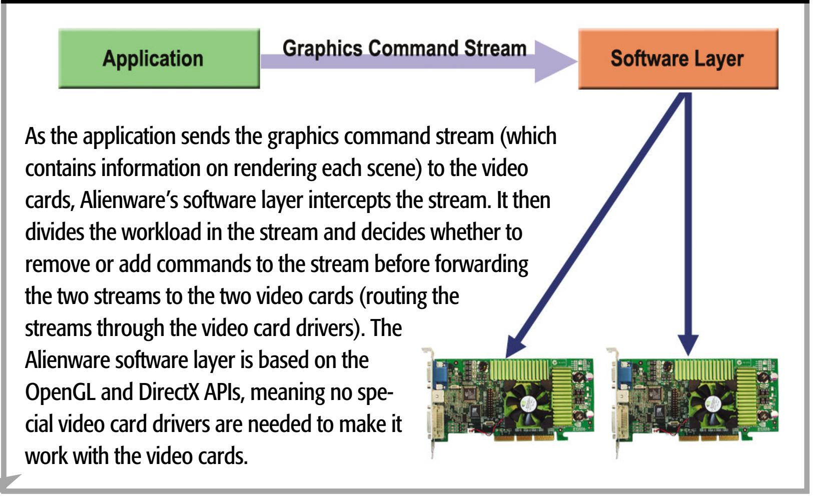

As the application sends the graphics command stream (which contains information on rendering each scene) to the video cards, Alienware's software layer intercepts the stream. It then divides the workload in the stream and decides whether to remove or add commands to the stream before forwarding the two streams to the two video cards (routing the streams through the video card drivers). The Alienware software layer is based on the OpenGL and DirectX APIs, meaning no special video card drivers are needed to make it work with the video cards.

Alienware Video Array Architecture

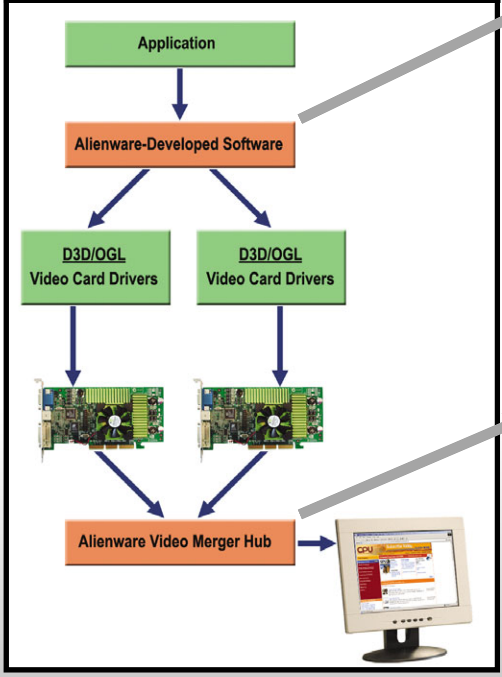

Alienware Video Merger Hub

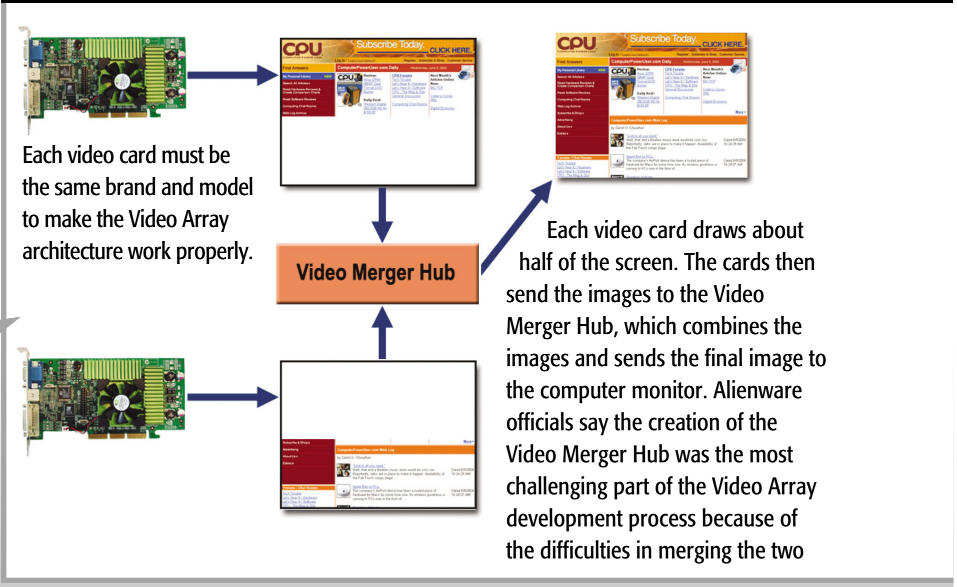

Each video card must be the same brand and model to make the Video Array architecture work properly.

Each video card draws about half of the screen. The cards then send the images to the Video Merger Hub, which combines the images and sends the final image to the computer monitor. Alienware officials say the creation of the Video Merger Hub was the most challenging part of the Video Array development process because of the difficulties in merging the two

Source: Alienware

Alienware does rely on some important newer technologies to make the overall ALX X2 system work.

PCI Express. As this key serial I/O technology finally makes it to the marketplace, the ALX X2 is poised to take full advantage of it. Because it's a serial bus and because it's scalable, PCI-E can achieve much higher bandwidths than previous technologies for moving graphics data, such as AGP. (You can read more about PCI-E on page 37 of the June 2003 issue of *CPU*.)

Organvidez says PCI-E is key to making Video Array and the X2 a reality, in large part because its scalability allows the dual-video card setup to work. "PCI Express was an essential technology, but only because we couldn't find a motherboard with multiple AGP ports on it, nor could we find a bridge chip that would allow for multiple AGP ports on a motherboard," he says. "This, along with the limitation of AGP 8X bandwidth, prevented us from being able to come to market with the Video Array solution any sooner."

Turnwater chipset. The ALX X2 motherboard will use the Turnwater chipset from Intel (also named the E7525). Turnwater is a high-end chipset that's designed for workstations and servers, which enhances the ALX X2 as a unique desktop option because most desktop computers use a consumer-level chipset. ALX X2 will use DDR2 memory in part because of the needs of Turnwater.

Xeon CPUs. Alienware chose Xeon CPUs because the high-end Xeons are SMP (shared memory multiprocessor)-capable, something Pentium 4 CPUs are not. SMP capabilities are important to help offset potential problems with CPU lag.

Customized case. Alienware has implemented several important features and technologies in the ALX X2 case, allowing it to deliver the power needed. For example, Alienware developers worked extensively in creating a unique liquid-cooling system that could handle the high-end cooling needs of a system with two GPUs and two CPUs, all of which can be factory overclocked by 10% or more. (Other available settings will allow customers to increase the overclocking even further, causing even greater cooling needs.)

"In order for the Alienware ALX line of systems to perform at peak performance and reliability, the custom liquid-cooling solution was designed to ensure [that] components remained within optimal operating temperatures, even when overclocked," Organvidez says.

Power At A Price

Alienware says the dual-video card ALX systems will achieve at least a 50% performance gain over single-video card systems. Depending on the other system components in place and the type of software running on the ALX, users will probably experience even larger performance gains. At the time of this writing, Alienware officials hadn't yet made specific, official claims to performance increases because they were still involved in the testing process. However, games and applications that are particularly graphics-intensive should receive the greatest performance benefits.

One potential problem: The dual-video card system might make graphics rendering so fast that the system will end up waiting for the CPU to catch up. Alienware officials say the X2's dual-CPU configuration using Xeons should alleviate most of the potential for CPU lag.

However, the biggest hurdle that the ALX X2 will have to overcome—the system's high price—is a result of its own success in creating the technology. By incorporating dual video cards, dual processors, liquid cooling, and other key technologies in the ALX X2, Alienware has created what some would call the ultimate gaming desktop computer. But including all of those components isn't cheap or easy, leading to an estimated price (at least $5,000 and possibly a few thousand more, depending on the customization features the customer selects) that some gamers on message boards began balking at immediately upon hearing the news of the ALX X2. Still, we're guessing there are plenty of dedicated gamers out there who will pay up.

by Kyle Schurman

ALX X2's Development

Alienware has been working on the idea for ALX X2 and Video Array for at least three years. Here's a snapshot of the development process.

2001. Developers at Alienware begin formulating the technologies required to build a dual-video card desktop PC.

2002. Alienware develops the Video Merger Hub hardware and an Open GL software solution that makes the dual-video card setup a possibility.

2003 (first half). Alienware showcases a sample dual-video card system to key companies in the graphics industry. As PCI-E technology begins to move toward implementation, Alienware developers begin tweaking their dual-video card solution to work with PCI-E.

2003 (second half). Alienware completes development of the X2 motherboard that allows Video Array technology to work.

2004 (first half). Alienware showcases an ALX X2 demonstration model at the E3 conference.

2004 (second half). Alienware will begin selling the ALX X2 desktop model.

images into a seamless final image. The video cards have to be fully synchronized to make the Video Merger Hub work properly.

The Video Merger Hub is the key to allowing Video Array to work at high resolutions because it handles the merging of the images and the sending of the combined image to the display without having to run the final image through one of the video cards again, a process that would greatly hinder the system's graphics processing power.

The Video Merger Hub connects to the computer through a USB connection, which allows the Alienware software layer to communicate with it directly, informing the Video Merger Hub about the graphics mode (such as 3D mode) and the workload ratio split.

LG FLATRON® LCDs.

Available in LG, XLG and XXLG.

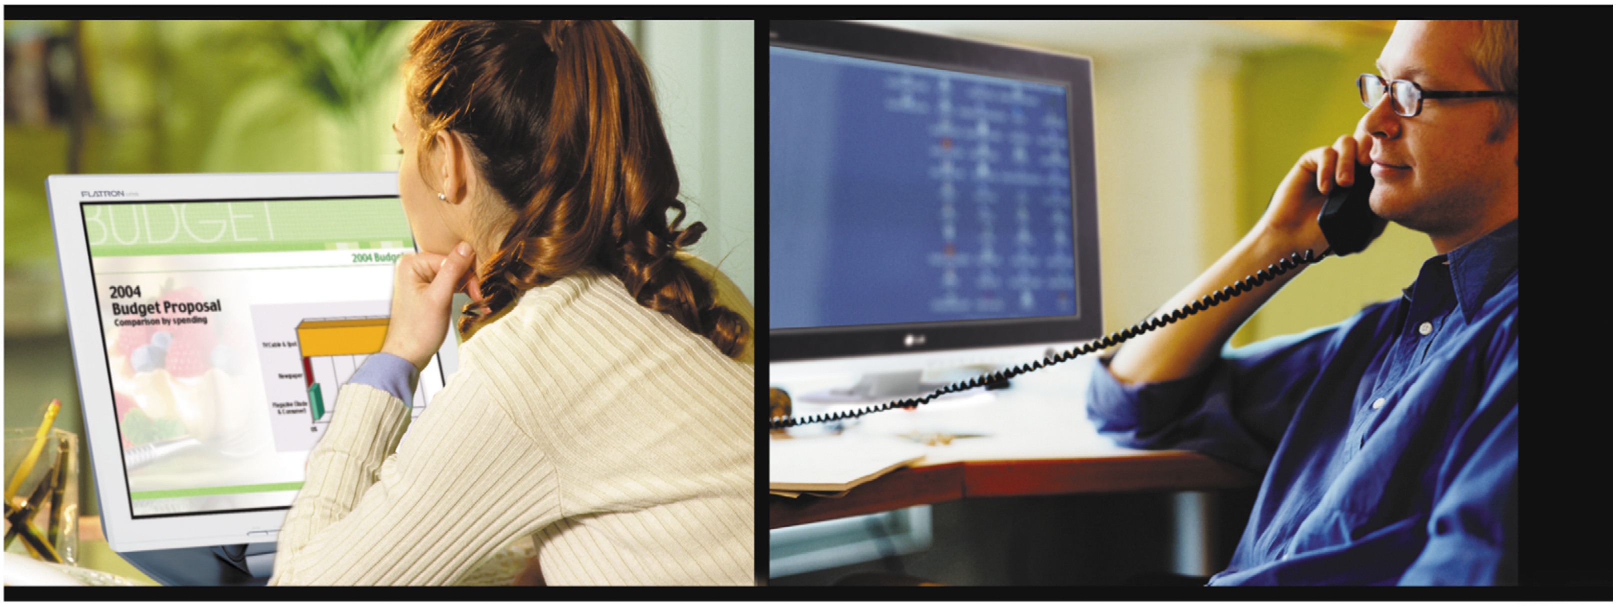

©2004 LG Electronics U.S.A., Inc.

L2320A - 23″ Wide-Screen HD/WUXGA LCD Monitor

Available at:

BEST BUY

BestBuy.com

Because one size does not fit all.

With screen sizes from 15″ to 30″, offering features such as DVI and HD inputs or TV function, there's a perfect LG fit for every home or office. LG LCD monitors offer exceptional quality, innovative features and are all backed by an outstanding 3-year replacement warranty.

The view is better with LG.

www.LGUSA.com

LG Life's Good

56 **Artificial Intelligence**
Boosting Your Brain With Software Assistants
by William Van Winkle

57 **Biometrics**
Your Body Will (Un)lock Your World
by William Van Winkle

58 **Biotechnology**
Can Technology Succeed Where Mother Nature Has Failed?
by Marco Chiappetta

60 **Cooling**
Dealing With Faster, More Powerful Processors
by Marco Chiappetta

61 **Grid Computing**
Not Just Folding@Home Anymore
by Dave Altavilla

62 **MEMS & Nanobots**
Ultra-Tiny Machines Will Reshape The World
by William Van Winkle

63 **OLED/Flexible Displays**
Tomorrow's Monitors Are Better, Brighter & Bendable
by William Van Winkle

64 **Optical Computing & Switching In Silicon**
Computing Power & Connectivity At Light Speed
by Dave Altavilla

65 **Quantum Computing**
Explaining The Inexplicable & Harnessing Its Power
by Dave Altavilla

66 **Robotics**
Lifelike Robotic Personal Servants Are No Longer Science Fiction
by Marco Chiappetta

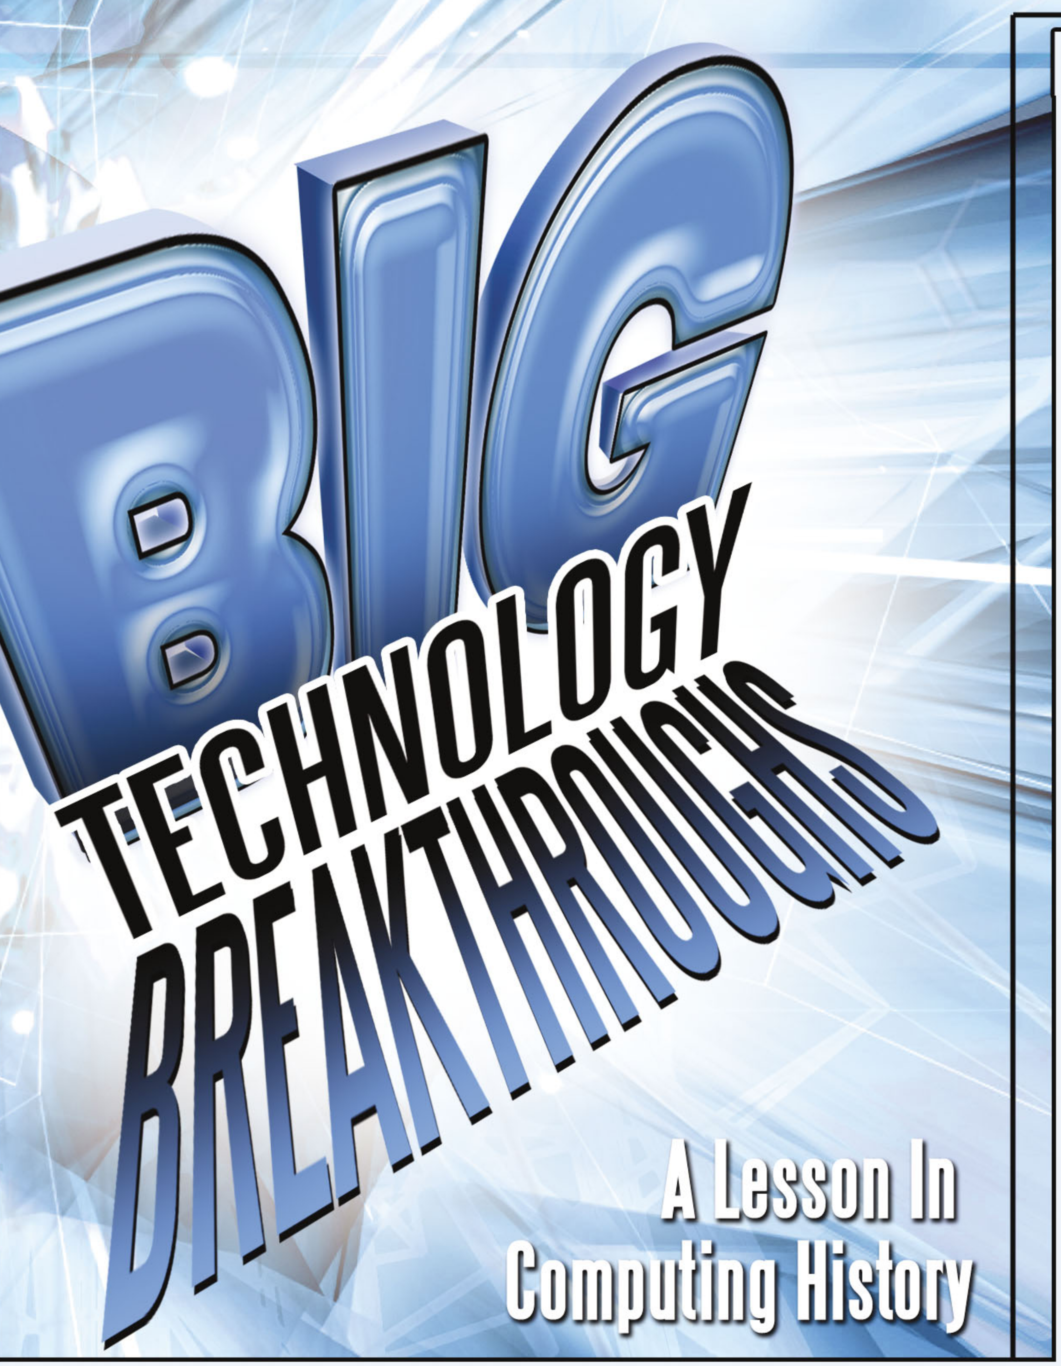

BIG TECHNOLOGY BREAKTHROUGHS

A Lesson In Computing History

As far back as 5,000 years ago, merchants used the abacus to calculate the cost of goods they sold. The device was refined over time, giving way to more efficient abaci. William Oughtred's slide rule followed in the 1620s, enabling faster multiplication and division operations using logarithmic functions. The slide rule enjoyed widespread use right up until the 1970s when calculators rendered them obsolete.

Computing went mechanical in 1642 when Blaise Pascal invented a numerical wheel calculator to help his father count taxes. Though the machine was limited to addition and subtraction, its implementation is mirrored in modern water meters and mechanical odometers.

The Computer Is Born

Up until the early 19th century, computing devices were generally limited to basic arithmetic calculations. Then, Charles Babbage conceived of a machine capable of calculating mathematic expressions using finite differences. He called it the Difference Engine and, in 1823, secured roughly $2,700 in grant money from the British government to build it. Unfortunately, manufacturing processes of the day weren't up to producing a number of identical parts, and after spending more than $30,000 on the project, Babbage had to halt construction.

Although the Difference Engine was mathematic in nature, it wasn't a general-purpose computer; it could only process numbers added in a specific sequence. Babbage's next project, the Analytical Engine, overcame that limitation, adding programmable input by way of a punch card reader. The machine was first conceptualized in 1837 and underwent continual revision until 1871, but it was never actually built due to financial and technical constraints. If it had been built, the Analytical Engine would have featured its own storage, with room for 1,000 numbers of 50 digits each; an arithmetic unit capable of addition, subtraction, multiplication, and division; and an output mechanism for printing and plotting curves. Estimations place the Analytical Engine at almost 100 feet long and more than 30 feet wide.

Because of its capacity for input, the Analytical Engine would have been the first computational device with a programming language, including provisions for loops and conditional branching. Lady Ada Lovelace wrote an algorithm for the machine in the 1840s, which would have allowed it to compute Bernoulli numbers. Consequently, she is often considered the first computer programmer.

IBM's Humble Beginning

Before Big Blue became the giant that it is today, the company's founder, Herman Hollerith, pioneered the use of electricity in computing with a punch card reader that helped analyze 1890 census data in the United States. The device read, counted, and sorted the cards, reducing a 10-year process to just one year of tabulation. It wasn't inspired by Babbage's work, either; rather, Hollerith's idea sprang from the system used by railroad conductors to read passengers' tickets.

Hollerith capitalized on his success by founding the Tabulating Machine Company in 1896, and his machines were subsequently used in several other countries for the same purpose. After making numerous revisions to the original design, including an automatic card-feeding mechanism and the first key punch, Hollerith's company became IBM (International Business Machines) Corporation in 1924.

The following decades saw punch card technology elevate in both public exposure and economic utility. Federal agencies began issuing punch card checks in the 1930s, and libraries started tracking books with punch cards by the '40s. But it wasn't until the 1970s that Herman Hollerith's contribution to computing gave way to more modern tools.

From Relays To Vacuum Tubes

The idea to use relays for computing originated in 1939 when George Stibitz built the Complex Number Calculator from telephone relays and crossbar switches, with three teletypewriters for input. In 1944, Howard Aiken and Grace Hooper designed the Mark I, which was similarly powered by electromagnetic relays. It was, however, much more advanced than Stibitz's early foray into relay experimentation. Prepunched paper tape fed the Mark I instructions to add, subtract, multiply, or divide. Programmed subroutines also facilitated the calculation of logarithms and trigonometric functions. It still involved an incredible number of mechanical parts for storing and counting data, though.

Soon after, relays began giving way to vacuum tubes, faster components that enabled what many consider to be the first generation of completely electronic computers. The ENIAC (electronic numerical integrator and computer) was perhaps the most notable example, unveiled in 1946 by J. Presper Eckert and John Mauchly. Used to solve atomic energy problems for the Manhattan Project, compute ballistic trajectories, and study cosmic rays, ENIAC was significantly faster than any preceding device. But all of its power came at a price. It involved more than 19,000 vacuum tubes, 1,500 relays, and an array of other components, consuming 200 kilowatts of power. ENIAC was divided into eight subsystems, with an IBM card reader and punch card charged with input and output, respectively.

Second-Generation Computing With Transistors

Late in 1947, two resourceful scientists working at Bell Labs developed the first point-contact transistor while searching for a solid-state replacement for the vacuum tube, which consumed too much power, ran too hot, and was far too unreliable. The initial implementation wasn't very practical, but a subsequent project called the junction transistor, clandestinely developed by team leader William Shockley, set the stage for an era of transistorized computing. Eight years passed before the transistor materialized in computers, though; the first application was telephony. Transistor radios weren't far behind, facilitating the success of Tsushin Kogyo, a small Japanese firm known today as Sony.

Bell Labs built the first fully transistorized computer, aptly named TRADIC (transistorized digital computer), in 1955. Equipped with nearly 800 point-contact transistors and 10,000 germanium crystal rectifiers, TRADIC consumed 100 watts of power and occupied 3 cubic feet of space. It wasn't quite as fast as the day's most potent machines, and transistors were substantially more expensive than vacuum tubes, but the tradeoffs in size, reliability, switching speed, and power consumption easily compensated.

TIMELINE

1837–Charles Babbage conceptualizes the Analytical Engine, revising it until 1871.

1884–Herman Hollerith applies to patent his punch card tabulating machine.

1939–George Stibitz builds the Complex Number Calculator at Bell Labs.

1946–The famed ENIAC is built by Eckert and Mauchly.

1955–The first digital computer, TRADIC, is built at Bell Labs.

1958–Jack Kilby of Texas Instruments demonstrates the first integrated circuit.

1969–Construction of the ARPAnet, predecessor to the Internet, commences.

1971–The Intel 4004, the world's first microprocessor, is unveiled.

1977–Apple is founded and introduces the Apple II.

1985–Microsoft launches Windows.

In the midst of the efforts being exerted to improve semiconductor technology, one arm of IBM's research team in San Jose was quietly developing the first magnetic disk drive, a device that would slowly mature, eventually setting the standard for high-capacity storage. Of course, that first drive, the 350 RAMAC (random access method of accounting and control) Disk File, was significantly different from today's refined Serial ATA devices. It consisted of 50 24-inch magnetic disks with enough space for 5 million 7-bit characters, about 4.4MB. The 350 RAMAC and its accompanying 305 RAMAC system, still based on vacuum tubes, sold for $190,000 dollars in 1956.

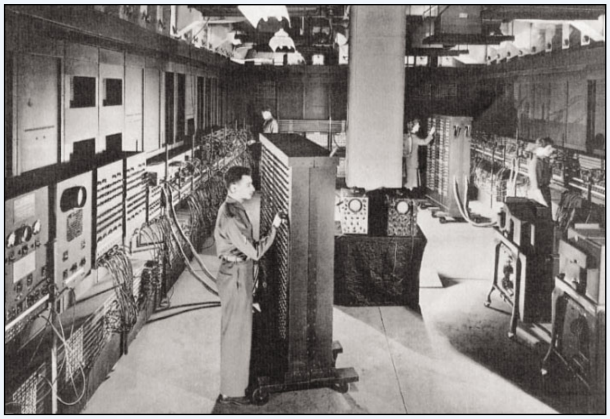
Arranged in a U-shape, ENIAC filled an entire air-conditioned room.

Defining Generation Number Three

Though the third generation of computers continued to employ transistors, they simultaneously shrank as more components were integrated onto single circuits. Jack Kilby of Texas Instruments was the first to demonstrate an IC (integrated circuit) in action, albeit a fairly simple one consisting of one transistor, three resistors, and a single capacitor.

Robert Noyce, co-founder of Fairchild Semiconductor and, later, co-founder of Intel, worked on a similar project, using silicon instead of germanium for the semiconductor material. Both Texas Instruments and Fairchild Semiconductor were awarded patents for their designs. The first commercially available ICs emerged in 1961 from Fairchild Semiconductor and were quickly adopted for use by the Air Force.

Three years later, Douglas Engelbart invented the first mouse designed for use within the context of a GUI, even if it wasn't officially unveiled until 1968 alongside Engelbart's NLS (Online System). Xerox's Bill English improved on that design in the 1970s, replacing perpendicular gear wheels with a single ball that could rotate in any direction.

Intel Introduces The Microprocessor

The development of ICs continued into 1971 when Intel unveiled the world's first microprocessor, the 4004. Miniaturization allowed Intel to integrate a number of functions onto a single piece of silicon measuring 3 x 4mm and consisting of 2,300 MOS (Metal Oxide Semiconductor) transistors. One year later, the 4004 was put to use on the Pioneer 10 spacecraft.

Meanwhile, IBM continued advancing the state of storage technology with its memory disk, an 8-inch plastic disk coated with magnetic iron oxide. The 5.25-inch floppy emerged in 1976 as a response to demand from Wang Laboratories for smaller disks. Sony launched the 3.5-inch drive in 1981, further refining the floppy to the point that it is still used in limited capacity today.

Two years after Intel unveiled its 4004, Robert Metcalfe of Xerox PARC explored the potential of a packet-based networking technology he called Ethernet. Defined primarily by wiring and signaling, Ethernet also incorporates protocol and packet format information, borrowing from the Internet Protocol suite. Metcalfe concedes that Ethernet actually evolved over the course of a few years. After receiving standardization in 1980, however, it quickly overtook the competing token ring and ARCNET technologies to become the most pervasive network topology right up until today.

The Dawn Of The Internet

Construction on ARPAnet commenced in 1969, but plans for the small, four-node computer network began seven years earlier when J.C.R. Licklider of the Department of Defense's ARPA (Advanced Research Projects Agency) first conceptualized general communication between the users of various computers.

The initial implementation of the network consisted of nodes at UCLA, Stanford Research Institute, UCSB, and the University of Utah connected to each other through 64Kbps digital lines. All four computers were DEC PDP-10s connected to secondary IMPs (Inter-face Message Processors), which served as routers. As other computers were slowly added to the network, compatibility issues emerged with the packet switching protocols used to communicate, necessitating a new architecture.

Robert Kahn and Vint Cerf worked together to shift the responsibility of reliability onto hosts rather than the network itself, minimizing potential protocol incompatibilities. Thus, in 1983, TCP/IP replaced NCP (Network Control Program) as ARPAnet's protocol. Military, educational, and research connections to ARPAnet expanded; commercial use was strictly prohibited.

Toward the end of the 1980s, the Department of Defense stopped funding ARPAnet. Private ISPs were formed around the same time to provide service to research networks and access to the public. By 1994, commercial providers had displaced the NSF (National Science Foundation) as the primary backbone to the Internet. Commercial restrictions subsequently disappeared, and the Internet as it's known today started taking shape.

Building Tomorrow's Technology Today

Technological advancements in the preceding century accelerated at breakneck velocities. The evolution of transistorized computing promises to uphold that trend, too, diverging into a bevy of exciting new disciplines that will eventually replace today's top-end gear.

The following pages will explore a representative sample of different fields enabled by today's technologies. Some of them, such as grid computing, biometrics, advanced cooling, robotics, and AI agents, are already commercially available. Optical switching and flexible LCD screens are also close at hand; experiments in each have yielded positive results. Other technologies, such as quantum computing and nanobots, may very well redefine the future but are still in the very early stages of development. CPU

AMD 64 Athlon FX

Introducing the AMD Athlon™ 64 FX processor. Take your system to extremes.
Double the data path from 32- to 64-bit and you more than double the thrill factor. Uninterrupted, ear-splitting, streaming audio and rich, razor sharp video make your pad a launching pad. What's more, you get all the power you need to edit, mix, and model your own digital creations with memory to spare. Prepare to blow minds. Get a dose of the AMD Athlon 64 FX edge at www.amd.com/amdathlon64fx

© 2003 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, AMD Athlon, and combinations thereof are trademarks of Advanced Micro Devices, Inc.

Artificial Intelligence Agents

Boosting Your Brain With Software Assistants

What

Since Alan Turing proposed in 1950 that a "thinking" computer would be one that provided query responses indistinguishable from that of a human, the field of AI (artificial intelligence) has endeavored to create computer systems that could behave like humans. Science fiction often envisions this as intelligent robots, although more realistic, achievable forms tend to be software programs (or parts of programs) designed to help with specific tasks.

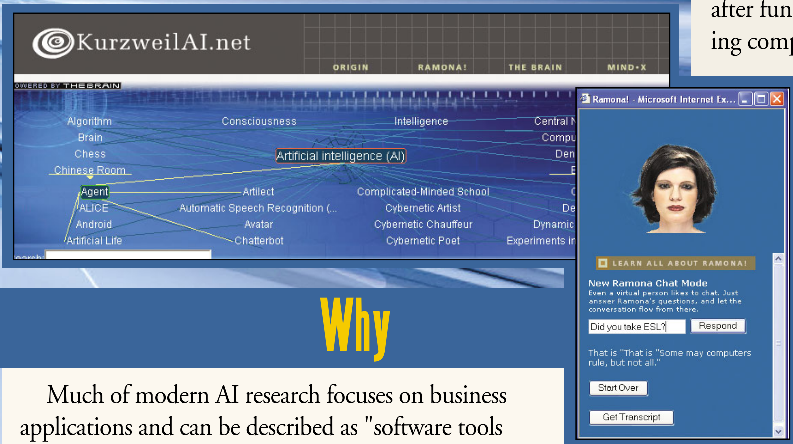

New Ramona Chat Mode
Even a virtual person likes to chat. Just answer Ramona's questions, and let the conversation flow from there.

Did you take ESL?

That is "That is "Some may computers rule, but not all."

Raymond Kurzweil is one of the most outspoken proponents of AI. His site contains substantial information across many AI disciplines, as well as a chance to chat with Ramona, his twitchy and hopefully still-under-development greeter agent.

Why

Much of modern AI research focuses on business applications and can be described as "software tools that make computing tasks better and easier." Those spelling/grammar underlines in your word processing docs are examples. Another app is natural language searching, such as that which Ask Jeeves popularized.

The Jeeves butler image represents a different AI field, intelligent agents, that some industry heavy hitters are now developing, including Microsoft and IBM. Agents are software programs (or parts of programs) that have enough intelligence to act on your behalf. A travel "agent" would know enough about your history and preferences to assemble multiple itineraries for a trip without you guiding it at each step. Another burgeoning field within AI R&D is search agents.

"There's no reason to think that search is only something that you do exclusively by stopping everything else you're doing and typing a query into a search input field," says Eric Horvitz, senior researcher and group manager, Adaptive Systems and Interaction Group at Microsoft Research. "There are many ways to have a system sort of know what you're looking at and automatically, in the background, make queries from that material and bring back relevant items and information for what you're working on at any given moment."

Horvitz's group is now looking at agents that help sort and present data and appointments. For example, tomorrow's agents will recognize tasks in email and either auto-schedule them or prompt the user for approval. The distinction Horvitz draws is that whereas computers now are tools, agents will soon enable computers to be collaborators.

How

There are many approaches to AI. The "knowledge-based system" holds that the foundation of intelligence relies on a basic understanding of a wealth of concepts. You cannot reach conclusions without facts to process. The 20-year-old Cyc project (www.cyc.com) encompasses thousands of common sense "rules," such as that when people die they stay dead. The system is programmed to deduce conclusions from these rules by way of an inference engine.

A more common AI technology is the neural network, modeled after functions of animal brains. Neural networks "learn" by drawing comparisons against a large database of examples. For instance, when a speech system has enough waveform data on various pronunciations of "dog," the system should be able to understand that word when nearly anybody speaks it. Neural networks are critical in fields such as aerospace and data mining.

Another form of AI hinges on fuzzy logic. Whereas Boolean logic deals with absolute truths, fuzzy logic deals with partial truths and can process them mathematically. Fuzzy AI makes its way into everything from meteorology to HVAC systems.

When

AI is present in some form in nearly every field of industry and is nearly ubiquitous throughout computing. Experts agree guesses of when a computer system will be truly intelligent or self-aware are nothing but blind fancy. As to when AI agents will arise, that depends on the level of sophistication one expects. If you expect a multiplatform agent to assist you with everything from finding a 5-year-old email to haggling on your behalf with creditors, you've got a long wait. Task-specific agents may be right around the corner, however.

"It wouldn't surprise me if people started relying more on automated, intelligent agents to assist them with tasks over the next five years," says Horvitz. "I don't think it'll happen all at once. It'll just be more and more functionality that you don't really notice happening."

Chances

No one disputes that AI will continue to grow and evolve. A 2003 report from Business Communications Company states that the total worldwide AI market in 2002 was $11.9 billion. BCC expects this market to grow over 12% annually and reach $21.2 billion by 2007.

Biometrics

Your Body Will (Un)lock Your World

What

Biometrics is the field of determining or verifying identity through the use of physical attributes, such as hand geometry, iris patterns, voice, or fingerprints. More obscure forms include analyzing vein patterns and behavioral biometrics such as walking gait, handwriting signature, and time spacing of typed keystrokes.

This is an image you won't see in any biometric database. When scanned, biometric images such as fingerprints are converted into abstracted numeric codes that can't be reversed back into their source image.

Modern fingerprint scanners, such as this U.are.U 4000 sensor from DigitalPersona, are conveniently small, affordable enough for most users, and far more foolproof than scanners of yesteryear.

Why

Biometrics replace passwords and PINs with the touch of a finger or glance into a camera. Aberdeen Group research has estimated that companies spend $100 to $350 annually per user setting up and maintaining passwords. Across a corporation, biometrics can save thousands in IS costs.

Ask the Department of Homeland Security, and you'll likely hear that widespread use of biometrics is essential for locating terrorists and keeping unwanted types off of U.S. soil. In a post-Sept. 11 world, biometrics for national security gets a lot of headlines if not a lot of results.

Biometrics in civil administration seems to be faring much better. Most European Union nations integrate fingerprint data into their ID cards, and the British government expects to deploy mandatory ID cards using fingerprint, iris, and facial recognition data by late 2005. A 2002 poll by the American Association of Motor Vehicle Administrators found that "87% [of 800 polled] favor making driver's licenses more resistant to tampering and counterfeiting by using biometrics."

In 2003, Harris Interactive released survey results showing that 13 million Americans had been the victims of identity theft or fraud and that their out-of-pocket expenses came to $1.5 billion annually. If identity profiles and financial accounts were encrypted with a data lock that only a biometric key could open, such fraud likely would shrink drastically.

How

Biometrics hinge upon creating mathematical abstractions based upon analysis of analog biometric input. Take fingerprinting, for example. The locations on the finger where epidermal ridges start and end are known as minutiae points. The software that DigitalPersona's U.are.U. 4000 Sensor uses identifies up to 70 different minutiae points, analyzes the relationships between the various points, and then uses this information to create a numeric profile of the fingerprint.

"We never memorize the image of a fingerprint," says DigitalPersona CEO Fabio Righi. "We create a mathematical representation of the fingerprint that we call a template. It averages about 300 bytes, and you cannot reconstruct the image of the fingerprint from that information."

The company's 4000 Sensor can determine if there is only air or a denser object, such as skin, on the window above the 4000's sensor. This ability helps thwart would-be fraudsters who would fog up a sensor window or shine a flashlight through it to have the scanner's CCD read the latent finger image the prior user left. Similar precautions to these also exist with other good biometric equipment.

When

From airport security to Hollywood blockbusters, biometrics are already embedded in our culture. At the consumer level, enthusiasts can already use desktop biometric devices for password management and to (un)lock encrypted files. Some financial institutions, particularly those in Europe, are starting to require biometric authentication for users to access their accounts. However, no one is yet guessing when biometric scanners and the wireless networking to link them will be so cheap and ubiquitous that they render the need for cash obsolete.

Chances

Last year, the International Biometrics Group forecasted that the global biometrics market would move from $719 million in 2003 to over $4.6 billion in 2008. For 2004, IBG shows fingerprint technology with a 48% market share, facial recognition 12%, hand geometry 11%, iris scanning 9%, voice printing 6%, and signature recognition 2%.

In mid-2003, NTT DoCoMo released its first biometric phone. Featuring a fingerprint scanner in its handset base, the Fujitsu F5051 sold 100,000 units in its first week. Different fingers can be registered to handle various tasks. NTT DoCoMo's eventual goal is to get people comfortable with biometrics so they can use their phones as secure wallets for m-commerce.

Biotechnology

Can Technology Succeed Where Mother Nature Has Failed?

What

Biotechnology encompasses the entire field of advances made in the biological sciences. There are numerous commercial biotechnology firms already in existence that produce genetically engineered substances for a variety of medical, agricultural, and ecological uses, and there are a multitude of firms using technology to improve, repair, or enhance the human body. Some of the most exciting developments in biotechnology are implants that may one day help the blind to see or help those with motor impairments to control a computer or electronic limbs. Amputees may even be able to control bionic limbs grafted to their bodies and nervous system, just like normal arms and legs.

Why

The focus of most biotech research has typically been to improve the quality of life of the general public or the physically disabled. The agricultural research has yielded some genetically engineered foods that grow faster, are more resilient, and have helped to feed the hungry, and electronic implants have been created to enable the blind to see or the paralyzed to operate a computer. A few of the most promising advances will help those with debilitating ailments regain their freedom and perform normal activities.

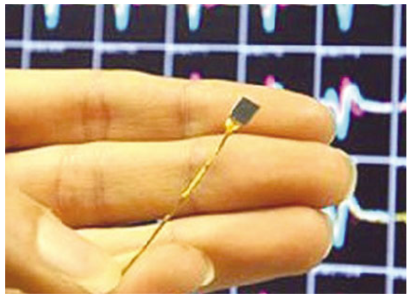

The BrainGate Neural Interface System, which is already in clinical trials, is intended to provide severely disabled people with a permanent, direct, and reliable interface to a personal computer.

How

The human body is essentially a group of systems all working together under the control of the brain. Because the brain, the entire nervous system, and the receptors in our muscles all use electrical impulses in some way, it's "only" a matter of interfacing with these systems and deciphering the electrical impulses to incorporate synthetic devices into the human body.

As long as the nerves and neural pathways into and out of the brain are still functional and intact, synthetic body parts and implants can be produced that will interface with the nervous system and the brain. Electrical impulses sent from the brain can then be used to control an electronic device, which in turn can send feedback signals back to the brain.

When

A few advanced devices already exist and are being tested in a limited number of humans. Cyberkinetics, for example, has already begun an FDA-approved clinical trial of its revolutionary BrainGate Neural Interface System, which "uses a proprietary technology for sensing the firing of hundreds of individual neurons within the cortex of the brain, transforming this information into a digital signal and transmitting this signal outside the body." Tim Surgenor, president and CEO of Cyberkinetics, said of BrainGate, "The ultimate goal of the BrainGate development program is to develop a safe, effective and unobtrusive neural interface which will allow quadriplegic people to control computers and computer-controlled devices using their thoughts." The company will be reporting its initial results from the trial by the beginning of 2005. In all probability, the trial will be a success, as monkeys implanted with an earlier version of the Brain-Gate were able to use it to play a simple video game.

Sandia National Laboratories has also developed a prototype MEMS-based implant that can be inserted onto the retina of a blind patient, with the hopes of being able to produce images good enough for the patient to read large print and to distinguish between objects. Sandia project leader Kurt Wessendorf said, "They won't be able to drive cars, at least in the near future, because instead of millions of pixels, they'll see approximately a thousand. The images will come a little slowly and appear yellow. But people who are blind will see."

Chances

The odds of devices like these becoming somewhat commonplace are quite good. Cyberkinetics already has plans to miniaturize the BrainGate and ultimately produce a wireless version; the current model requires a wire that runs from the patient's skull to the computer interface. Sandia's part has already been successfully tested. As the technology matures and we begin to better understand the inner workings of the human brain, devices like these should only get better. Implants may even perform better than our biological designs one day. High-resolution eye implants may enable us to see objects way off in the distance like a set of binoculars, devices like the BrainGate may allow us to control computers with only a thought, and bionic limbs may give us the ability to lift hundreds or even thousands of pounds. The technology is already being discussed, and early prototypes are currently in existence.

Give *your* car the luxury of satellite navigation.

© 2003 Garmin Ltd. or its subsidiaries

The StreetPilot® III from Garmin has features found in expensive in-dash navigation systems, plus one very important quality the others don't have. StreetPilot III is portable, so you can have the best in GPS navigation in any car you drive. Look up addresses, attractions and other services, and StreetPilot III will create a route and provide turn-by-turn directions with voice prompts to get you there.

StreetPilot III and GPS V from Garmin. Navigation for every car.

GARMIN. www.garmin.com NASDAQ GRMN

Available at: BEST BUY COMPUSA

Cooling

Dealing With Faster, More Powerful Processors

What

The microprocessors used in today's personal computers are composed of millions of transistors and run at clock speeds approaching 4GHz. Unfortunately, because they operate at such high speeds and incorporate so many transistors into a tiny space no bigger than a thumbnail, they also generate a significant amount of heat. More complex graphics processing units run at lower clock speeds, but they too are made up of millions of transistors and pump out quite a bit of heat, as well. To function reliably, these chips need to be cooled to prevent permanent damage and to keep them operating within specifications. The most popular ways to cool microprocessors today are with heatsinks and fans, watercoolers, and in rarer cases, with vapor-phase (refrigeration) coolers.

Why

Excessive heat will damage any electronic component, so it is extremely important to channel heat away from the device. Traditional cooling techniques have served us well in the past, but the processors of tomorrow could generate so much heat that these customary cooling methods may no longer be practical. Computers, cellular phones, and PDAs are also getting smaller, which makes cooling the processors within them much more difficult. The need to keep future processors running reliably, without using extremely loud fans or prohibitively expensive techniques, has fueled some advanced research into alternative cooling methods. Some significant progress has been made in the development of nanofans and in the use of nanolightning to create ion-wind currents.

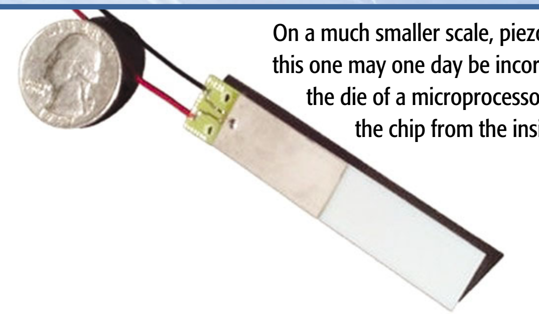

On a much smaller scale, piezo fans like this one may one day be incorporated into the die of a microprocessor, cooling the chip from the inside.

How

Nanofans are essentially tiny piezoelectric fans that flap back and forth to create small gusts of wind. Because they have no gears or bearings and have only one moving part, piezofans are quiet and reliable and generate very little heat of their own. And because they don't use magnets, piezofans don't create any electromagnetic noise that can interfere with delicate electronic circuits. A team of engineers at Purdue University has been working to perfect the piezofan and plan to shrink them down to about 100 microns. Piezofans that small will fit right onto a processor and compliment the traditional coolers already in use.

Two members of the Purdue University Cooling Technology research project, Daniel J. Schlitz and Vishal Singhal, have also recently founded a new company called Thorrn Micro Technologies to pursue commercial applications for a technique that uses ion-winds to cool a processor's core. By applying an electrical current to two opposing electrodes, ions flow from one electrode to the other, like miniscule lighting bolts. In the process, the ions collide with neutral air atoms and cause them to move in the same direction of the ion flow. This reaction is known as the "corona wind effect" or an "ion-wind." By using tiny, strategically placed carbon nanotubes between the electrodes, Thorrn Micro Technologies was able to direct heated air atoms through a chip's core and expel them from the side.

When

Piezofans have existed for years. It will, however, take some time to find the right materials and to shrink them to the small sizes necessary to incorporate them into microchips. Experts estimate the technology is three to five years away from being used in any commercial applications. Microsized, ion-wind-driven devices, however, are still in the lab research phase, but the technology has shown real promise. Despite the relative immaturity of the devices, Daniel J. Schlitz has stated that they could be ready for commercial applications in as little as two years. Considering he is one of the founders of Thorrn Micro Technologies, though, Schlitz's estimates may be a bit optimistic. Products such as the "Ionic Breeze," available from The Sharper Image, already employ this technology, albeit on a much larger scale, to silently move air without the use of fans.

Chances

The future of these technologies is not clear. With advances in the design and manufacture of microprocessors, the need for these devices may be diminished. Current leakage is a major reason why today's microprocessors run so hot. If we find new materials or techniques to drastically reduce the amount of current leakage, and the efficiency of current designs is improved, the useful lifespan of traditional cooling devices may be increased. There are also a number of other advanced cooling techniques being explored. Electron Evaporation was recently developed out of the need to cool the world's largest electromagnet. In the end, the technology that works best, and that costs the least to produce, will probably prevail. What we do know is that a simple heatsink and fan alone won't be enough.

Grid Computing

Not Just Folding At Home Anymore

What

If ever there were a next “big idea” to follow the all-powerful Internet, grid computing could very possibly be it. Although grid (or distributed) computing applications initially started out in the research labs of such places as distributed.net and UC Berkeley’s SETI (Search for Extraterrestrial Intelligence; seti.ssl.berkeley.edu) project, it has potential for much broader market appeal in the corporate sector and perhaps even in the consumer sector.

Grid computing provides its users with the ability to use system resources such as computational power, bandwidth, and storage across disparate organizations and geographies, distributing workloads, maximizing IT infrastructure, and allowing for new capabilities in collaboration. That’s a mouthful for sure, but it’s not a completely new idea to even many casual computer users these days. In short, think of a “grid” of computational resources all acting as one huge virtual processing, transmit, and storage element. Companies such as IBM, HP, Sun, and Microsoft are now mobilizing an effort to bring this new “big idea” to market.

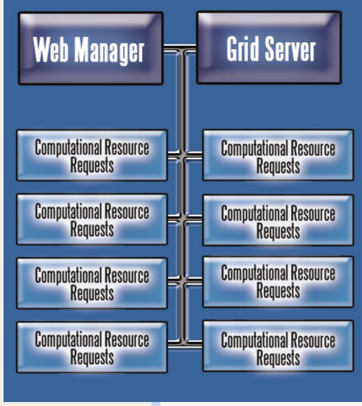
Web Manager
Grid Server
Computational Resource Requests
Computational Resource Requests
Computational Resource Requests
Computational Resource Requests
Computational Resource Requests
Computational Resource Requests
Computational Resource Requests
Computational Resource Requests

Web Manager and Grid Server systems manage computational resource requests and workload distribution across a grid of processing server resources.

Why

The folks at Berkeley were challenged with processing the huge amount of data that was generated from their radio telescope, which was listening for stray artificial signals possibly coming from distant stars. As you can imagine, the amount of data coming in from the telescope and the level at which they needed to analyze it drove the need for the most powerful supercomputers to handle the workload. Even if they had the funds for this costly process, they still didn’t have enough horsepower to get the job done. So the team developed a client that people could download and run on their local machines, which would then request data sets to be processed from SETI’s server and upload the data back to the server when the data set was complete. On a larger scale, corporate entities have needs similar to the SETI workload, where large data sets need to be processed for a myriad of applications.

How

Hardware and software are the two components that make up the basics for grid computing. At the hardware level, there are typically two or three types of required systems: a Web server, a grid manager/server, and engines. These elements make up the grid that a client can connect to and use as a resource. In high-end architectures, the grid manager and Web server are separate elements, as well as the engines, which are the computational or storage resources.

More simplistic architectures can have a combo Web/grid manager system on the front end that is both providing an interface to the grid and managing workload distribution among the engines. Engines can be standalone servers, blade server architectures, or even workstations (those with higher levels of computational and storage power) across the network. SETI’s implementation used end-user PCs that were essentially turned into an engine resource across SETI’s basic grid network.

Taking things one step further is the concept of utility computing for the average corporate or home user. Users simply run thin-client machines that do not actually have software licenses, high-end processors, or large amounts of storage. Instead, users lease computational resources and application usage from the central “utility” server and engine farm in their area for a monthly fee.

When

Grid computing is here today, and it’s here to stay. Analyst firm IDC states that the worldwide grid computing market is expected to exceed $12 billion in revenue by 2007, across high-performance computing technical markets and commercial enterprise. Obviously, large system OEMs such as IBM, Sun, HP, and even smaller grid/utility hopefuls such as Egenera (www.egenera.com) are gunning for this. In addition, if Uncle Bill Microsoft has anything to say about it, grid or utility computing will come to the end-user masses in the not-so-distant future, as well. We’ll believe that when we see it, but five to 10 years out, when broadband is the standard rather than just the majority, anything is possible.

Chances

You could almost bet the farm on grid computing becoming very pervasive in the enterprise in the next three to five years. Even Intel’s efforts focused around blade server architectures give credence to the ever-changing demands on the data center. The real $64,000 question is: Will it make it to the mainstream end user looking to run Word and Excel and start up a Quake 4 Death Match? Technology and infrastructure to support this kind of evolution for the home user is still out on the roadmap horizon somewhat.

MEMS & Nanobots

Ultra-Tiny Machines Will Reshape The World

What

A MEMS (micro-electromechanical systems) is a class of devices that builds very small electrical and mechanical components on a silicon chip. A single MEMS is generally measured in microns, although multiple MEMS can be packaged into a larger device measuring up to several centimeters. MEMS technology first reached the market in the 1990s in the automotive industry, such as in airbag sensors. MEMS are now pervasive in everything from inkjet cartridges to toxic chemical detectors.

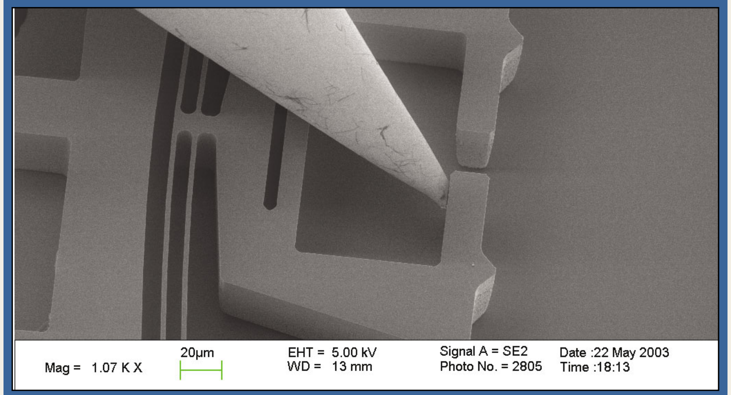

No, this isn't a fractal image. This Zyvex rotary actuator device is used in MEMS fabrication, perhaps as a gear that will help move a motor. Its teeth measure 3 microns in width, which enables the motor to take a 6-micron step.

Logically, if macroscale electromechanical components can make such things as power tools or a data storage system, it follows that microscale components should be able to make the same sort of things, only smaller. This has led some experts to ponder the advent of "nanobots," robots manufactured at the nanoscale performing such tasks as environmental waste cleansing or cellular DNA repair.

Why

In many cases, MEMS technology is more cost effective. In the airbag example, older systems used fairly large accelerometers in the vehicle's front linked to electronics next to the airbags, with the whole subsystem costing over $50 per vehicle. In a MEMS implementation, everything on one chip is $5 to $10.

Equipped with wireless communication, MEMS or nanoscale systems could be built into building materials and alert people to developing structural problems. They could cover miles of lunar landscape in solar cells to power a colony. The question isn't why we want such devices; it's when can we build them?

How

Modern MEMS are made much like CPUs. You begin with a substrate of silicon or other material; add additional material layers; and etch out features, such as beams, mirrors, or hinges. MEMS are simple devices, like an axe compared to a chainsaw. Moreover, they're still relatively large. Making complex tools at the nanoscale requires nanoassembly.

There are two basic approaches to nanoassembly: top down and bottom up. In essence, the top-down approach holds that it's easier to make a 1mm device with a 1cm tool than a 1-meter tool. You build tools that can make smaller tools that can make smaller tools and so on, with each generation of tools working in parallel until you can build a device at the desired scale. The bottom-up approach seeks to build devices from bare molecules or atoms. Nature does this all the time, but researchers are struggling to find a way to do this that doesn't require countless labor hours of manually moving particles with electron probes or laser tweezers.

When

"We're about two and a half years into our top-down program," says James Von Ehr II, CEO of Zyvex, a leading company in MEMS and nanoassembly research. "We're hopefully going to launch this bottom-up program soon if the government funds it, and the odds look good that it will. Within five years we ought to have the top-down finished and the bottom-up well in hand, enough to merge those two over another five-year period in order to get massively parallel, atomically precise manufacturing. That is still a long way from a nanoassembler or people's wildest dreams of nanotechnology, but it would clearly be the first steps."

Ultimately, the chief hurdle that exists in ultra-small MEMS and nanobots may be physics. Essentially, quantum effects limit device size. Magnetic fields or stray radio waves may scramble functionality. At some point, it's possible that dust motes may pose a physical threat to MEMS. We are still probably at least a decade from such worries, though.

Chances

In July 2002, In-Stat/MDR issued a report stating, "Every major market has now embraced the technology." However, in some segments, "such as the communications and consumer markets, MEMS has only just begun to scratch the surface. As a result, worldwide revenues for MEMS are forecast to grow from $3.9 billion in 2001 to $9.6 billion in 2006."

The National Science Foundation issued funding numbers last year showing that the U.S. government was keeping pace with or exceeding Western Europe and Japan on nanotech funding. However, the projected U.S. budget for 2004 was $849 million, which seems low considering the potential economic benefits. Of this, $5 million comes from the EPA and $222 million from the DoD. For more, see the National Nanotechnology Initiative site at www.nano.gov.

OLED/Flexible Displays

Tomorrow's Monitors: Better, Brighter & Bendable

What

OLED (organic light-emitting diode) displays use organic materials that, when properly charged, give off a bright glow—significantly brighter and sharper than LCDs. While glass has been the substrate of choice for most OLED displays, attention is shifting to plastic (particularly thin-film plastic manufactured in long rolls) as a means to facilitate flexible displays.

Why

LCDs are at an inherent disadvantage because they rely on power-hungry backlights that are then shuttered by the liquid-crystal material and polarizing filters used in TFT monitors. Because OLED pixels are electroluminescent, meaning they give off their own light when charged, there's no need for a backlight or filter between the pixels and the viewer. This is why OLED displays are considerably brighter and sharper than conventional flat-panel displays.

Moreover, OLEDs conserve far more power. A glowing OLED pixel requires between 2 volts to 10 volts but only needs this power when active. A black OLED pixel requires no power. In contrast, TFT displays must power their backlights constantly and apply power to LCD cells to block that light from escaping.

OLED pixel-response rates are more than 100 times faster than LCD, according to Kodak (which holds many of the original OLED patents), thus making OLED a better fit for movies and gaming. Kodak claims its organic materials deliver a wider color range than LCD, far better luminance, and wider viewing angles (170 degrees) and can be implemented in screens a mere 1.5mm thick.

PHOTO COURTESY OF PHILIPS

This image of a Philips 13-inch PolyLED TV prototype gives you an idea of the image quality an OLED display is capable of.

How

OLEDs use thin layers of organic material sandwiched between a transparent anode and a metallic cathode, all of which is usually built onto a glass or plastic substrate. With proper voltage, the organic materials produce light. The specific structuring of the layers, the materials in the anode and cathode, and the fluorescent molecules used to "dope" the electroluminescent layers all contribute to the efficiency and precision of OLED color output. The active-matrix addressing often used to control TFT pixel cells in today's screens also carries over to OLED displays.

There are two general types of organic materials in OLED displays: "small molecule" and "macro polymer." Most OLEDs thus far have used small molecule materials, but the technology encounters serious challenges as screen size increases, including thermal warping of the shadow mask.

Royal Philips Electronics is betting that polymer displays are the direction of the future, both in reflective and emissive varieties. Reflective-polymer technology is typified in Philips' collaboration with E Ink, the result being a pigment-based "epaper." (See April 2004's "Under Development," page 107 for more.) Emissive polymer OLEDs, though, have their electroluminescent molecules suspended in a paint-like liquid that can be literally printed through inkjet nozzles.

When

OLED screens have already appeared in devices such as Kodak's EasyShare LS633 digital camera and MSI's MEGA Player 515. The first 15-inch (1,280 x 720) OLED prototype, designed by Kodak and Sanyo, went on display in October 2002. Three years earlier, state-of-the-art OLED prototypes had been 2.5-inch displays.

Some experts believe roll-to-roll OLED display manufacturing should quickly become less expensive than LCD technology once it starts volume production. Some believe this may happen in 2005, although this seems a bit optimistic.

Chances

According to market analysis firm iSuppli, the 2003 global OLED panel market was just $250 million (17.3 million units). Many were small panels that went into cellular phones and other portable electronic devices. However, iSuppli expects the OLED market will reach $470 million this year (36.2 million units), up to $3 billion by 2007, and possibly exceed $4 billion in 2010 (more than 366 million units).

Small-molecule OLED technology has a head start, but manufacturing depends on a vacuum chamber and other costly elements. If makers can "print" polymer displays, particularly on flexible substrates, the market could change, with large displays being made relatively cheaply.

Optical Computing & Switching In Silicon

Computing Power & Connectivity At Light Speed

What

Optical computing, otherwise known as photonic computing, is the process of using light waves and pulses instead of electrons to transmit data. It uses an on/off state of light rather than the traditional on/off state of electricity to represent binary 1s and 0s used in computing and communications systems around the world. In the networking space, big iron communications equipment providers such as Lucent, Cisco, and Nortel have been designing and building equipment based on this technology for years. Fiber-optic communications power the backbone of the Internet, as well as your phone line and many other public and private networks. However, only until very recently have researchers been able to prove its viability in local (as opposed to across a network backbone) computing and switching applications.

Why

The need for another means of transferring data was obvious to networking communications equipment providers. Electrons were no longer cutting the mustard, and optics were a natural migration, as the bandwidth limitation is at the speed of light, which is orders of magnitudes faster than the speed of electrons. In addition, technologies such as DWDM (Dense Wavelength Division Multiplexing) allow networks to host multiple connections over a single fiber, making it possible to transmit voice, data, video, and audio over a single fiber. Bandwidth makes the world go 'round in networking, and the same holds true for modern computing.

As processors begin to scale from 3GHz to 10GHz in the next five years or so, their connection speeds at the system level are also going to have to scale, as well, in order to make efficient use of their computational power. For example, a 10GHz processor with an 800MHz system bus is a rather large technological mismatch. After all, what good is that 600 horsepower Porsche engine if you can't wind it out on the straightaways?

Intel's breakthrough modulator, built in standard silicon, can encode data onto a laser beam at gigahertz speeds.

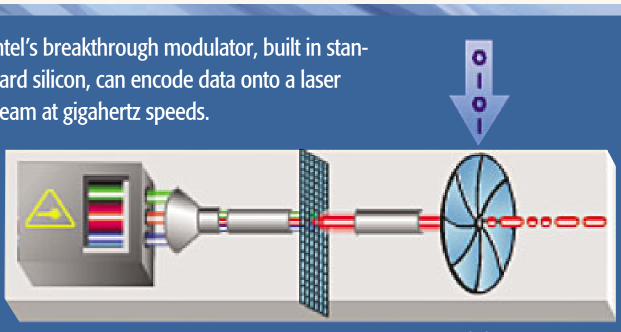

PHOTO COURTESY OF INTEL.

How

The catch to making optical computing a viable solution for personal computing connectivity is cost. Until recently, optical communications technologies could only be implemented in high-end expensive III-V semiconductor compounds such as Gallium Arsenide (GaAs). Then along come companies such as Intel, who are now pioneering silicon photonics. However, researchers haven't been able to prove out silicon photonic switching speeds (on/off signals that represent binary 1s and 0s) in excess of about 20MHz. The bottleneck has been the fiber-optic modulator that is used to encode data on the ultraviolet light wavelengths that are used in transmission. Intel now has found a way to encode the data on the light beam at 1GHz and has done it purely in silicon with standard manufacturing technologies and die geometries that are available today. These new silicon-based optical modulators can now be used to build high-speed optical communication links at the system I/O level, enabling an unprecedented system and port level bandwidth.

You can also envision this technology at the chip level, breathing new life into and extending Moore's Law. Intel feels it'll be able to scale this technology up to 10GHz and beyond in the future. Want to know what's beyond PCI Express? Think optical interconnects. They are faster, have no problem with signal-to-noise degradation, and are immune to crosstalk and EMI (electromagnetic interference).

When

Optical computing could very well pave the way for bringing an information and computation superhighway direct to future PC, workstation, and server designs near you in the next five to 10 years. This technology is a natural migration from relatively slow microelectronic processing and communications to computation and communication at the speed of light. Pat Gelsinger, senior vice president and chief technology officer at Intel, notes, "It is the kind of breakthrough that ripples across an industry over time, enabling other new devices and applications. It could help make the Internet run faster, build much faster high-performance computers, and enable high-bandwidth applications like ultra-high-definition displays or vision-recognition systems."

Chances

Not long ago, highly skilled engineering candidates coming out into the workforce with electrical engineering degrees were being scooped up like gold dust. But the new frontier is likely to be micro-optics and silicon photonics. Better bone up on your photonics because it looks like electronics could even someday go the way of the dinosaur. Also, if you're placing your bets on Silicon Valley continuing its comeback, we'd say it's just a matter of when, but the future still looks as sandy as a beach.

Quantum Computing

Explaining The Inexplicable & Harnessing Its Power

What

Get ready for this one because it's a doozy. Do whatever you do to relax and open your mind a bit: meditation, breathing exercises, whatever works for you. You're going to need it. In order to understand the concept of quantum computing (or computing relative to the unique phenomena that is observed in quantum or "nonclassical" physics), you must first understand quantum mechanics a bit, or at least what scientists have learned about this new frontier.

Remember, relax and focus, people. . . . Take, for example, what has widely been illustrated in the area of quantum mechanics as the Two Slits experiment. It's pretty simple on the surface. Shine a light source through two parallel slits on a piece of cardboard or similar material and observe the result on a screen. You should see two parallel lines on the screen in front of you. Simple right? Not quite. What you'll actually see is a diffraction pattern that looks like several bars of light across the screen. Scientists have proven that a single photon (or particle of light) can cause interference with itself and seemingly travel through both slits at once and show up on the receptor surface twice! Good to see you, Neo, welcome to the "real" world.

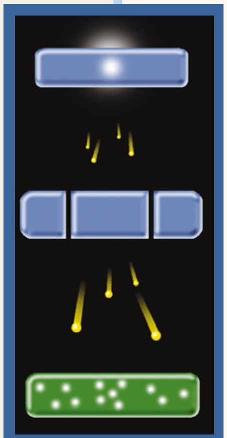

Waves of photons causing interference with themselves create a classic refraction pattern on the screen in the Two Slits experiment.

Why

This experiment is a proven fact of observation, only scientists and philosophers are still very much trying to understand why it happens, and there are many theories on the topic. Also, this is only one simple example of quantum mechanics and its associated phenomena. There are many others, and at the atomic level, matter in general also obeys the laws of quantum mechanics, which are very different from the laws of physics that control electronic microarchitectures such as logic gates.

The start of assembling some mathematical sense to this "freak of nature" of sorts is its fundamental building block called a qubit or quantum bit. Basic computer chips operate at the bit level, represented by a pattern of 1s and 0s. The "on" state represents a 1, and the "off" state represents 0. An atom can also assume a composition of two different electronic states. As a result, this two-state representation called a qubit can represent either a 1 or 0 and have at any given point in time a value of either 1 or 0. Pretty deep, right? You bet, and this concept speaks to the very essence of reality in the "physical" world. Is the signal on or off? Actually it's both on and off. We know, beam me up, Scotty, my brain hurts.

How

Now in theory, if we can develop a computational algorithm that makes use of quantum technology, we'll have exponentially more powerful code engines that will solve exponentially more complex problems in exponentially less time. What's more interesting is that a scientist by the name of Peter Shor, from AT&T's Bell Laboratories in New Jersey, developed the first quantum algorithm that supposedly can perform efficient factorization. Brute force prime factorization is the technique that code breakers must use to decrypt security keys, such as those generated from RSA algorithms for security in applications worldwide. Shor's algorithms supposedly have potential to decrypt RSA keys with relative ease if a capable quantum computer can be built, rather than running the equations on traditional computational architecture. Shor states, "Quantum computing won't just crack codes; quantum cryptography will protect future systems." He goes on to say, "The first quantum computer is still very much a thing of the future, but at least now scientists and physicists are working toward building one with renewed interest."

When

There are lots of smart guys like Peter Shor who dare to imagine, so you never know when quantum computing may actually become feasible—next month, next year, or maybe not even in your lifetime. The possibility exists, though, because scientists continue to think beyond the physical domain and into the quantum domain.

> "The first quantum computer is still very much a thing of the future."

Chances

You can think of quantum computing as possibly the new frontier, sort of like back in the day when Buck Rogers was firing laser guns and people thought of it as science fiction. We work with lasers now in many areas of technology, and they're certainly no longer sci-fi. Science and research can very possibly bring fiction into reality.

Robotics

Lifelike Robotic Personal Servants Are No Longer Science Fiction

What

Robots have been used on assembly lines for quite some time, but these mechanical workhorses seem like toys when compared to the humanoid robots being developed by some talented engineers from across the globe. There are currently about 70 major anthropomorphic robotic projects in existence in numerous countries, with the most advanced designs coming out of Japan and the United States. There are many smaller projects in the works, as well. Most of these robots are being designed to look, act, and move like humans. They'll eventually be used for space travel, protection, assistance, and entertainment, among other things.

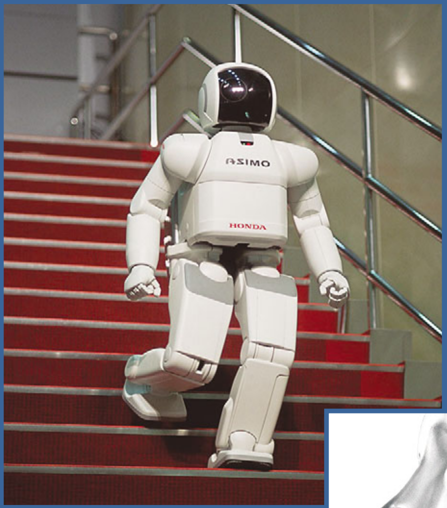

Honda's ASIMO (Advanced Step in Innovative Mobility) is an advanced anthropomorphic robot that can walk, climb stairs, and perform a number of simple tasks just like a human.

The NS-5 is the star of this summer's "i, ROBOT," featuring Will Smith and Bridget Moynahan.

Why

The idea of a robotic personal servant has been around for decades. Science fiction writers have imagined robots of all kinds, ranging from the quirky R2D2 to the overly analytical Lieutenant Commander Data. Each one was built to serve a purpose, whether it was to act as an electronic wingman or pursue a more humanlike existence. The robots of today, and the immediate future, however, don't have such lofty goals. Most are being designed to assist people with mundane everyday activities or to facilitate the lives of the physically disabled. Having a personal servant to assist paralyzed patients, for example, will give them much more freedom and may eliminate some of the difficulties associated with their daily care. Some anthropomorphic robots are also being designed to protect humans in dangerous situations, such as diffusing bombs or cleaning areas contaminated by deadly gases or biological agents.

How

To date, the talented engineers dedicated to designing anthropomorphic robots have taken a very methodical approach. Each individual part and subsystem has been developed independently, fine-tuned, and ultimately brought together to produce a completed humanoid robot. It took years to develop the legs, motors, servos, and gyros that work well enough to allow a robot to walk adequately. Arms, wrists, and fingers were also made to mimic the movements of humans. Visual and auditory systems needed to be produced so the robots could be somewhat aware of their surrounding environment. And all of these different subsystems, among others, had to be integrated together and coalesced with a myriad of sensors and microprocessors to make a single, humanlike robot. The software to control all of the individual subsystems had to be created, as well.

When

Lifelike humanoid robots already exist in multiple forms. NASA, Honda, Fujitsu, Sony, and many other entities, along with numerous private engineers, have working prototypes in their labs at this very moment. The problem with the existing designs is that they lack the maturity and refinement necessary to be considered useful. They're also very expensive to produce. Honda's ASIMO, for example, can climb stairs, walk, and even pick up objects, but dropping millions of dollars on a prototype to watch it perform these simple tasks hardly seems like a worthwhile investment. The good news is that virtually all of the major advances in the field have taken place in the last 18 to 20 years, which is a relatively short time. If the technology continues to advance at its current pace, lifelike and very capable robotic assistants may be available in the next decade or two.

Chances

With the major technological advances made in the field of advanced anthropomorphic robots, they will inevitably be made available to the general public. Some of the more sophisticated models can already walk, communicate, and recognize their owners. As they become stronger and more agile, get faster processors and more memory/storage, and incorporate more powerful artificial intelligence algorithms, these humanoid robots will emerge as capable assistants and perhaps even become loyal companions. The technology will surely evolve, and with the inescapable reduction in the costs of the materials used to produce robotic assistants, they will probably be relatively affordable one day, too. If you long for your own version of C3PO, be patient: It's not a matter of "if" but "when."

MAXIMUM PERFORMANCE

NOMINAL PRICE

Experience the new 8X DVD+/- RW Burner from AOpen

Now take advantage of the superior performance of AOpen's® new 8X DVD +/-RW drive at an amazing price!

AOpen – the world's leading PC component manufacturer – has pulled-out all the stops with this new feature-loaded, dual-format DVD burner. It not only writes both DVD+R and DVR-R media at true 8X speeds, it also serves as a handy, fast and convenient CD media writer and player. It may be a record-breaker, but there's one thing the DWR8800 DVD Burner is NOT, and that's a budget-breaker.

With extra goodies like Ahead's *Nero Burning ROM, InCD, VideoStudio 7,* CyberLink's *Power DVD* and 321 Studios' *DVD X Copy* (trial version) – you'll make movies, author professional-caliber DVDs and efficiently manage all your media storage needs – all with one nifty package.

It's fast. It's reliable. It's affordable. It's the AOpen DWR8800 8X DVD Burner. Maximum performance. Nominal price.

AOpen®

www.aopen.com
1-888-97-AOpen

© 2004 AOpen America Inc. All rights reserved. AOpen is a trademark and the AOpen logo is a registered trademark of AOpen Inc.
All other trademarks are the property of their respective owners.

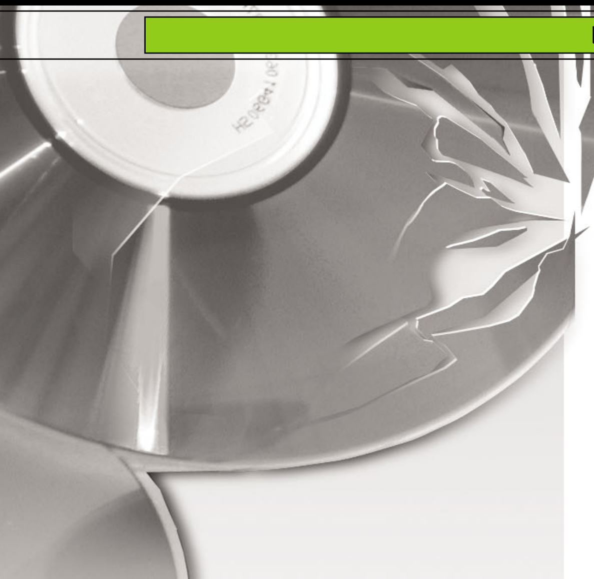

by Warren Ernst

The Bleeding Edge Of Software

Inside The World Of Betas

Hymn .61 beta

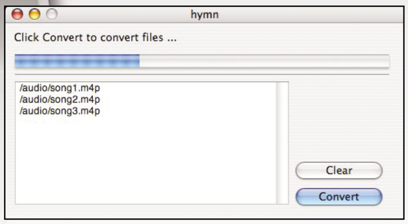

Click Convert to convert files ...

/audio/song1.m4p
/audio/song2.m4p
/audio/song3.m4p

Official product name: Hymn
Version # previewed: .61 beta
Publisher: The Hymn Project
Developer and URL: The Hymn Project; hymn-project.org
ETA: Q3 2004
Why you should care: MP3 users can finally buy music at the iTunes Store.

Apple's iTunes Music Store is a runaway hit, but if your digital music player isn't an iPod, you're out in the cold. The service's song format is M4P, and songs are encoded with the AAC copyright-protection scheme, which prevents you from copying them. It also prevents playing the songs with software other than iTunes or on MP3-only hardware. Though there's no AAC/M4P-to-MP3 converter (we know of), the Hymn beta software can remove the AAC protection.

Hymn is a command-line program for Windows/Mac OS X/Linux, but the Mac version has a drag-and-drop GUI. Documentation is sparse, but for the command-line versions, typing **hymn /?** lists all the command flags. Here's a Windows tip: Drag Hymn to the Desktop and then drag your M4P files from Windows Explorer to the Hymn icon. The songs will convert to standard M4A files quickly. This only works one file at a time.

All the platforms require that iTunes be installed, and on the Mac you also need an iPod connected. The converted files still have copyright holders' names (along with your name) encoded within the file, so sharing the files on the Web would lead authorities back to you. Hymn isn't about music theft; it's about playing music on your software/hardware. Apps like GX::Transcoder can convert the M4A file to an MP3 to play on MP3-compatible hardware/software.

A lack of a Windows GUI is really the main indicator that Hymn is a beta product. Otherwise, it worked flawlessly on dozens of songs. If you've been shying away from the iTunes Music Store because of the file format, Hymn may be your answer. ▲

The Poster Toaster 1.0.48 Beta

You wouldn't create an entire Web site in Notepad, so why run an eBay business with a similarly simplistic tool? If the time-consuming online forms that eBay requires you to complete put you off, check out The Poster Toaster.

Many eBay auction posters require enrolling for a paid service or paying a lot for what should be an inexpensive app. TPT has the competition beat here; it's free and doesn't require any account beyond eBay and the basics your ISP provides. If you have access to an FTP server, TPT has everything it needs to add auctions, manage and post images, and send email to buyers.

Setting up TPT can be vexing, however. You have to configure your servers and FTP settings manually. A PDF help file and a Test These Settings button helps, but you need a little knowledge of the Internet to interpret any error messages.

TPT's UI also takes some getting used to, but adding auctions is surprisingly fast. You'll organize auctions by folders on the left and enter auction properties at the right. You can enter HTML with a nice embedded WYSIWYG HTML editor and create or use templates with the same editor. TPT uploads your photos to your FTP site, posts the auction on eBay, and verifies the links.

Although there are warnings regarding its beta status, TPT was stable and bug-free. If you can slog through the setup, TPT will definitely save you time. ▲

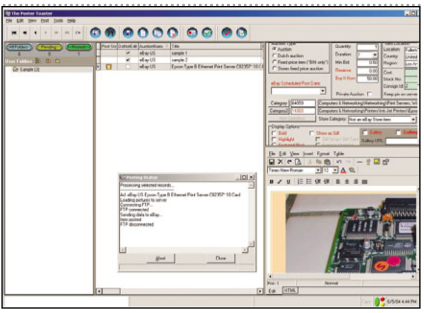

Official product name: The Poster Toaster
Version # previewed: 1.0.48 Beta
Publisher: The Poster Toaster
Developer and URL: The Poster Toaster; www.theposterto aster.com
ETA: Q3 2004
Why you should care: A cheap, easy way to auction a ton of stuff on eBay.

Shooting Star Beta 2.5.7

If you think the Poster Toaster is a good idea, you'll love Shooting Star. The Poster Toaster only creates and posts auctions, but Shooting Star does pretty much everything else you need, including track dates, payments, inventory, and email communications. Though the program's user interface is a bit clunky, the program's depth and breadth is astounding.

After plugging in your eBay and PayPal login information, plus a mailing address and basic ISP mail server information, SS logs into eBay and retrieves your buying and selling auction information. Double-clicking an item displays the auction data, such as start and end times, current prices, number of bids, and more.

The program's real power comes when you communicate with buyers. When an auction finishes, SS can calculate shipping and handling charges and send those calculations via email to the buyer. When you receive a payment, you can log it in SS, which pops up another email form to inform the buyer that shipment is imminent. Optionally, SS prints a packing sheet and shipping label and tracks dates. If a payment is late, a color-coded system displays this and generates an email to the buyer. You can also easily create email templates, but those that SS bundles are professional and friendly.

SS also reminds you to post feedback about buyers and sellers. In addition, it can email them with feedback reminders. The app also summarizes and generates reports on auctions over time, including profit/loss, item reports, the number of bids received, auctions pending, and more.

At $50, the application's price is a little steep, but eBay junkies swear by the program. The 30-day trial period is a good time to determine your needs, but try the program when you have a few auctions going at once. ▲

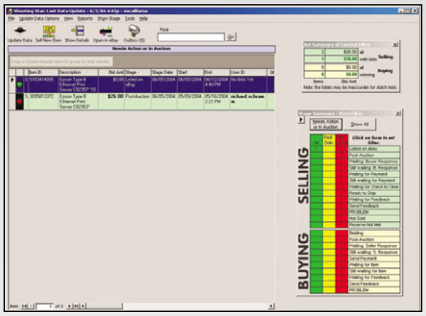

Official product name: Shooting Star
Version # previewed: 2.5.7 Beta
Publisher: Foo Dog Software
Developer and URL: Foo Dog Software; www.foodogsoftware.com
ETA: Q3 2004
Why you should care: Shooting Star may be the easiest way to handle dozens of auctions.

McAfee CleanBoot 1.0 Beta 3

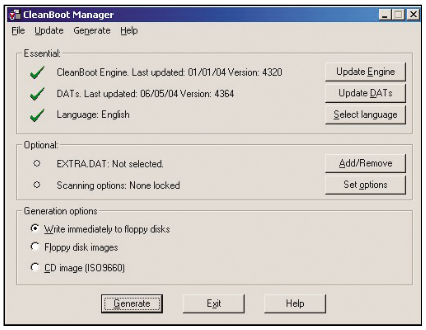

Essential:
✓ CleanBoot Engine. Last updated: 01/01/04 Version: 4320
✓ DATs. Last updated: 06/05/04 Version: 4364
✓ Language: English
Optional:
○ EXTRA.DAT: Not selected.
○ Scanning options: None locked
Generation options
● Write immediately to floppy disks
○ Floppy disk images
○ CD image (ISO9660)

Official product name: CleanBoot
Version # previewed: 1.0 Beta 3
Publisher: Network Associates
Developer and URL: Network Associates; us.mcafee.com
ETA: Q4 2004
Why you should care: Potentially an IT pro's very good friend.

Killing viruses is becoming a full-time job for IT pros, but when an infection is so severe a machine can't boot, there aren't normally a lot of options. Usually, you must pull the infected hard drive out and place it into another machine with antivirus software, which obviously is inconvenient. But what do you do when there isn't a second machine readily available? Use McAfee CleanBoot—at least, when it gets released.

CleanBoot builds a bootable version of McAfee VirusScan via a bootable CD ISO image or floppies. If you create images, you must create the actual discs, which means if you have a bootable USB flash "keychain" drive, you can probably create something.

Once created, the boot process runs a DOS-like VirusScan in a DOS environment, reading/writing to the NTFS partitions native to Windows 2000/XP/2003. Obviously, FAT16 and FAT32 partitions are compatible. The latest version recognizes 91,110 viruses, Trojans, and variants and will clean files where possible, delete when necessary, and scan the Registry. Once it runs through the system, it should clean enough to boot on its own steam, where more thorough cleaning can happen. It also identifies infections by name, so you can download the fix-it programs you may need.

The beta was rough around the edges; the CD image it created wouldn't boot when burned. You must download virus definitions manually and then manually point CleanBoot to the files. Generating discs or images under Windows 98 was hit and miss. However, the floppies it created worked fine, and CleanBoot located some recently installed viruses from my virus zoo. If McAfee polishes the package, CleanBoot may end up in IT pros' bag of tricks. ▲

Send Us Your Betas: Know of software in the beta stage that's deserving of some attention? Let us know. We'll take a look at it and possibly give it a go-round. Send your prospects to bleedingedge@cpumag.com.

UP TO SPEED

Upgrades That'll Keep You Humming Along

Media players are at the forefront this month with new releases of Microsoft's Media Player, WinDVD, and J. River Media Center. Yahoo! has also been busy with a new messenger and toolbar with a spyware killer that tries to one-up Google's.

Updates

GameCam 1.1

GameCam lets you record gameplay audio and video while playing. This version improves support for ATI and NVIDIA cards and overall video-capture performance. It also fixes issues with unresponsive hot keys and gamma correction in OpenGL titles.

Get it at: www.planetgamecam.com

J. River Media Center 10

This full-featured player aims to be a one-stop multimedia repository by managing music, video, digital photos, MP3 players, and even personal video recordings. This version adds a new UI with improved media file browsing, tagging, device synching, DVD burning, and 3D visualizations.

Get it at: www.musicex.com

LimeWire 4.0.4

Changes in the P2P file-sharing client for the Gnutella network include a new UI with more convenient button placement and more comprehensive search feedback. LimeWire also scans more listening ports at startup and reduces intrusive error messages from the P2P network while making other error messages more descriptive. Numerous rare bugs were also fixed.

Get it at: www.limewire.com

Microsoft Office 2004 for Mac

This majro upgrade for Mac OS X users improves Excel charts and adds smart buttons to Word and Excel. The Project Center hub lets you access all Office data, including email, documents, and contacts.

Get it at: www.microsoft.com

Microsoft Windows Media Player 10 Technical Beta

The public test for the next WMP has a revamped UI, better support for external devices, an ecommerce engine for purchasing media, and a jukebox that integrates music, video, DVD playback, and recorded TV.

Get it at: www.microsoft.com

Opera 7.50

The alternative browser gets a new UI with better skinning and toolbar customization. A new integrated IRC supports smileys, buddy lists, and file transfers. The improved email client has faster indexing and learning modes for the spam filter.

Get it at: www.opera.com

The Bat! 2.11

This alternative email client gets email chat and chat folders for monitoring multiple chat sessions. The Quick Templates feature creates templates out of blocked text.

Get it at: www.ritlabs.com

WinDVD Platinum 6

InterVideo's DVD player software now includes support for such hi-def formats as DVD-A, WMV-HD, and Intel High Definition Audio. You can now run Real, QuickTime, and MPEG-4 within the player. New video-processing technology reduces jumpiness in fast-moving images, and the formerly separate Audio Booster Pack now is included for playback of Dolby Headphone, DTS Digital Surround Sound, and SRS TruSurroundXT.

Get it at: www.intervideo.com

Yahoo! Messenger

This update seems designed to look more like AOL's IM client, with integrated radio channels and a game engine so you can hear tunes or play against chat buddies. Users get avatars and sound effects to accompany responses, easy image sharing, and client skins.

Get it at: www.yahoo.com

Yahoo! Toolbar Beta with Anti-Spy

This newest version adds an icon that lets you scan and remove known adware, spyware, system hijackers, hostile scripts, some worms and Trojans, and keyboard recorders. The tool alerts you when new versions are available with updated spyware profiles.

Get it at: www.yahoo.com

Driver Bay

ATI Catalyst 4.6

The new release for all Radeon cards lets users optimize the appearance and positioning of HDTV output and adds OpenGL acceleration for ATI Rotation mode. Pixel Shader 2.0 and OpenGL refinements increase performance for some games, and the drivers improve performance under ATI's new X800 line of cards.

Get it at: www.ati.com

Creative WebCam Notebook Driver 1.04.01

This driver lets Photoshop directly access the camera's images in Windows XP and provides a camera icon for the Scanner and Camera Wizard.

Get it at: www.creative.com

nForce ForceWare Unified Drivers 4.24

For nForce and nForce2/3 motherboards, these WHQL certified drivers add multiple audio upgrades, including a new mixer, a speaker setup tool, virtual surround sound for headphones, and support for Microsoft's new hi-def WMV-HD audio. The drivers also add SATA controller support, better RAID functionality, and better compatibility with CD rippers.

Get it at: www.nvidia.com

by Steve Smith

Broadband Phone Services

The Telephone Meets The Internet

Online addicts used to want an extra phone line to access the Internet. Today, the Internet can replace your telephone entirely. With a broadband phone service, you can make and receive calls using a phone without any assistance from your Baby Bell. Instead, the phone plugs into your Internet connection.

You need two things to turn your broadband connection into a personal phone line: a subscription to an Internet telephone service and a compatible phone. You can use a special SIP (Session Initiation Protocol) phone or regular phone plugged into a VoIP (Voice over IP) gateway. A SIP phone resembles a traditional phone but has a USB connector to plug into your computer or a 10Base-T connector to plug into your router. A VoIP gateway connects between your router and standard phone. This is convenient if you want to use a cordless phone or your light-up Hello Kitty model. Some gateways have two ports to which you might connect a phone and fax machine. Some services claim to work with fax; others don't guarantee success. (Faxing worked with all the services I tested.)

You Make The Call

Once connected, broadband phone services work like a regular phone: Place a call, and you'll hear a dial tone. When someone wants to call you, she can use any phone to dial your number, and your phone rings. A houseguest who wants to use your phone probably won't even know it's not a POTS (Plain Old Telephone Service) line. The only caveat is that some services require dialing 11 digits (1+area code+phone number) even if the destination phone number is in the same area code.

Each service I tested assigned a phone number for incoming calls, letting me choose the phone number's area code and prefix. Before choosing a service, make sure it offers phone numbers in the area you want. The number doesn't necessarily have to be in your physical area code, though; you can get a New York or San Francisco number, for example, which rings to your broadband phone wherever it is.

With the services I tested, phone numbers were tied to the SIP phone or VoIP gateway, which you'll need to purchase. You can plug the hardware into any broadband Internet connection and make/receive calls within a few minutes. Thus, you can take your phone (and phone number) with you.

The quality of calls for the seven services I tested ranged from fair to excellent. The best calls were indistinguishable from a good local connection. Services with imperfect voice quality usually sounded like a digital mobile phone, with slight glitches that didn't significantly detract from the conversation. A few services occasionally suffered from latency, which can lead to awkward pauses and both parties speaking at the same time.

The services were loaded with features that old-school phone companies charge extra for, such as call waiting, caller ID, three-way calling, call forwarding, and voicemail. Some offer features your phone company can't match, such as listening to voicemail on the Web and logging online calls you've placed, answered, and missed.

If you replace your traditional phone line with a broadband phone, use a service that supports 911. When you call 911 from a traditional phone, the phone company's Enhanced 911 service sends your address to a dispatcher's computer screen. Broadband phone services work differently. Some send a hard-coded address (which you provide the company) to the dispatcher. If you move the phone to another location and dial 911, an incorrect address will display. Other services don't support Enhanced 911. You may be connected to a local dispatcher, but he won't automatically know your location. If you're choking on a chicken bone, this could prove fatal. Also consider that traditional phone lines work when the power is out, but broadband services require power and a stable Internet connection.

For this roundup I tested using the VoIP gateway the respective service supplied. I connected each gateway to a cable modem through a home firewall/router. (The firewall was never an issue; I was able to make/receive calls with each service without opening any ports.) I made several U.S. calls and one to Germany with each service. While on the phone, I simultaneously used the broadband connection for other tasks, such as downloading large files and videoconferencing. These other online activities never noticeably affected call quality.

Chatting At The Computer

The services in this roundup rely on SIP (Session Initiation Protocol), a signaling protocol for Internet telephony, conferencing, and similar applications. (See www.faqs.org/rfcs/rfc2543.html for more on SIP.) SIP was developed in 1999, but Internet telephone services have really made many Internet users aware of it.

Even before SIP, Internet users were enjoying free long distance voice communication via programs such as Skype (www.skype.com) and AOL Instant Messenger (www.aim.com). The SIP services reviewed here aren't free, but their ability to connect to anyone with a standard phone is an important advantage over computer-based chat software. The people you talk with don't need a broadband connection or even a computer.

SIP-based services won't kill PC chat software. After all, many of the software choices are free and some include videoconferencing. But SIP phones will take toll-free Internet-enabled chatting from the realm of geeks to anyone who can use a phone. ▲

AT&T CallVantage

Maybe some traditional long distance providers are worried about Internet competition, but others are jumping onboard. AT&T's CallVantage has features and drawbacks that stem from AT&T's position as a major long-distance provider.

U.S. calls were perfectly clear with no noticeable latency. I suspect that's because calls spend a relatively short time bouncing around the Web before shifting to AT&T's traditional long-distance network. My call to Germany was a different story. My German friend called it the worst connection of the seven services, complaining of static when I wasn't speaking. His voice was perfectly clear to me, but when he wasn't talking, there was complete silence; I couldn't hear any background noise, which was disconcerting. This was one of two services, however, in which my caller ID information was received in Germany.

Setup was a hassle. Every other service let me plug the VoIP gateway into my Linksys router, but CallVantage would only work when the VoIP gateway was connected directly to the cable modem. I had to rewire the network, move my firewall rules to the VoIP gateway, and reconfigure the router. My FTP client also stopped working in this configuration. This might not be much of a problem if your VoIP gateway will stay put, but it makes it difficult or impossible to take your phone number with you on a trip.

CallVantage does have oodles of extras and offers local phone numbers in a large number of markets. At press time, phone numbers were available in 34 markets in nine states. By July, the service should be available in 71 markets in 21 states. You can also transfer your existing phone number to CallVantage, shutting down your POTS line in exchange for broadband.

Features include caller ID, call waiting, call forwarding, group conferencing with up to nine people, Locate Me for incoming calls, an online phone book, voicemail, and the ability to listen to voicemail on the phone or the Web. Unlimited nationwide calling is $34.99 a month, plus federal and state taxes. (Because AT&T is a traditional, taxed phone company, CallVantage does include taxes where some other services don't.) With a one-year contract, the $30 service activation fee and $60 VoIP adapter were included at no extra cost, a promotion that may still be available as you read this.

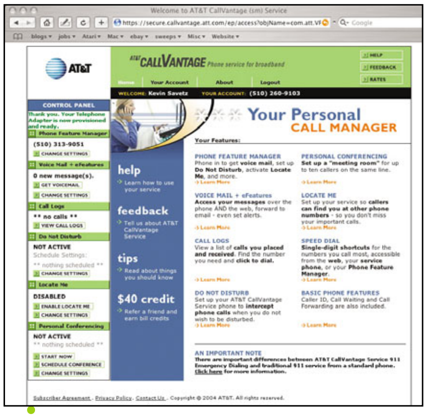

WELCOME: Kevin Savetz YOUR ACCOUNT: (510) 260-9103

Thank you. Your Telephone Adapter is now provisioned and ready.

(510) 313-9051

0 new message(s).

** no calls **

NOT ACTIVE
Schedule Settings:
** nothing scheduled **

DISABLED

NOT ACTIVE
** nothing scheduled **

Learn how to use your service

Tell us about AT&T CallVantage Service

Read about things you should know

Refer a friend and earn bill credits

Your Features:

PHONE FEATURE MANAGER
Phone in to get **voice mail**, set up **Do Not Disturb**, activate **Locate Me**, and more.

PERSONAL CONFERENCING
Set up a "meeting room" for up to ten callers on the same line.

VOICE MAIL + eFeatures
Access your messages over the phone AND the web, forward to email - even set alerts.

LOCATE ME
Set up your service so **callers can find you at other phone numbers** - so you don't miss your important calls.

CALL LOGS
View a list of **calls you placed and received**. Find the number you need and **click to dial.**

SPEED DIAL
Single-digit shortcuts for the numbers you call most, accessible from the **web**, your **service phone**, or your **Phone Feature Manager**.

DO NOT DISTURB
Set up your AT&T CallVantage Service phone to **intercept phone calls** when you do not wish to be disturbed.

BASIC PHONE FEATURES
Caller ID, Call Waiting and Call Forwarding are also included.

AN IMPORTANT NOTE
There are important differences between AT&T CallVantage Service 911 Emergency Dialing and traditional 911 service from a standard phone. Click here for more information.

. Copyright © 2004 AT&T. All rights reserved.

CallVantage
$34.99
AT&T
www.usa.att.com/callvantage
CPU CPU CPU

Broadvox Direct

With a broad array of extras and online account management, Broadvox seemed to have the makings of a contender, until nagging problems cropped up.

Voice quality was generally good but with noticeable latency with some calls. On one call I placed, there was no dial tone, but a later attempt let me place the call. My German associate and I agreed that the overseas quality was very good.

Receiving calls was more troublesome. Broadvox lets you set up phone numbers in different cities, which it calls Friends and Family numbers. When I set up a secondary number, it never worked. Callers to that number heard, "The subscriber you have dialed is not in service." Even calls to the primary number were dicey, with apparent unreliable routing. Calls sometimes ended up with a fast busy signal, fax tones, or someone else entirely. Sometimes my unanswered phone would continue to ring even after the caller hung up.

Service plans range from $29.95 per month for unlimited U.S. and Canadian calls to $12.95 per month for 500 minutes, with additional usage at 2.9 cents per minute. In addition to the usual features—voicemail, call waiting, caller ID, call forwarding, etc.—Broadvox includes Find Me, Follow Me, online call logs, and the ability to track missed calls on the Web. This all sounds good, but if you can't reliably receive incoming calls, the features are irrelevant.

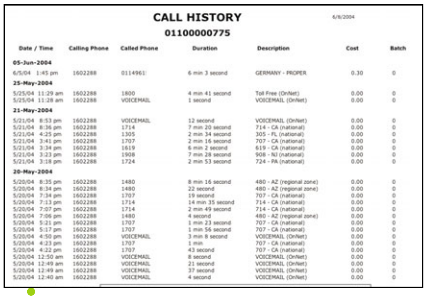

CALL HISTORY
01100000775
6/6/2004

Date / Time	Calling Phone	Called Phone	Duration	Description	Cost	Batch
05-Jun-2004						
6/5/04 1:45 pm	1602288	0114961	6 min 3 second	GERMANY - PROPER	0.30	0
25-May-2004						
5/25/04 11:29 am	1602288	1800	4 min 41 second	Toll Free (OnNet)	0.00	0
5/25/04 11:28 am	1602288	VOICEMAIL	1 second	VOICEMAIL (OnNet)	0.00	0
21-May-2004						
5/21/04 8:53 pm	1602288	VOICEMAIL	12 second	VOICEMAIL (OnNet)	0.00	0
5/21/04 8:36 pm	1602288	1714	7 min 20 second	714 - CA (national)	0.00	0
5/21/04 4:25 pm	1602288	1305	2 min 34 second	305 - FL (national)	0.00	0
5/21/04 3:41 pm	1602288	1707	2 min 16 second	707 - CA (national)	0.00	0
5/21/04 3:34 pm	1602288	1619	6 min 2 second	619 - CA (national)	0.00	0
5/21/04 3:23 pm	1602288	1908	7 min 28 second	908 - NJ (national)	0.00	0
5/21/04 3:18 pm	1602288	1724	2 min 53 second	724 - PA (national)	0.00	0
20-May-2004						
5/20/04 8:35 pm	1602288	1480	8 min 16 second	480 - AZ (regional zone)	0.00	0
5/20/04 8:34 pm	1602288	1480	22 second	480 - AZ (regional zone)	0.00	0
5/20/04 7:34 pm	1602288	1707	19 second	707 - CA (national)	0.00	0
5/20/04 7:13 pm	1602288	1714	14 min 35 second	714 - CA (national)	0.00	0
5/20/04 7:07 pm	1602288	1714	2 min 49 second	714 - CA (national)	0.00	0
5/20/04 7:06 pm	1602288	1480	4 second	480 - AZ (regional zone)	0.00	0
5/20/04 5:21 pm	1602288	1707	1 min 23 second	707 - CA (national)	0.00	0
5/20/04 5:17 pm	1602288	1707	1 min 56 second	707 - CA (national)	0.00	0
5/20/04 4:50 pm	1602288	VOICEMAIL	3 min 8 second	VOICEMAIL (OnNet)	0.00	0
5/20/04 4:23 pm	1602288	1707	1 min	707 - CA (national)	0.00	0
5/20/04 4:22 pm	1602288	1707	43 second	707 - CA (national)	0.00	0
5/20/04 12:50 am	1602288	VOICEMAIL	8 second	VOICEMAIL (OnNet)	0.00	0
5/20/04 12:49 am	1602288	VOICEMAIL	21 second	VOICEMAIL (OnNet)	0.00	0
5/20/04 12:49 am	1602288	VOICEMAIL	37 second	VOICEMAIL (OnNet)	0.00	0
5/20/04 12:40 am	1602288	VOICEMAIL	4 second	VOICEMAIL (OnNet)	0.00	0

Broadvox Direct
$29.95
Broadvox
www.broadvoxdirect.com
CPU CPU

iConnectHere Broadband Phone

While the VoIP gateways for all other services arrived preconfigured (or configured themselves the first time I connected

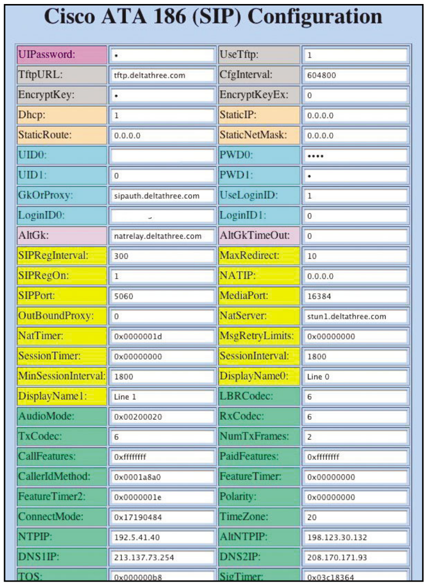

Cisco ATA 186 (SIP) Configuration

UIPassword:	•	UseTftp:	1
TftpURL:	tftp.deltathree.com	CfgInterval:	604800
EncryptKey:	•	EncryptKeyEx:	0
Dhcp:	1	StaticIP:	0.0.0.0
StaticRoute:	0.0.0.0	StaticNetMask:	0.0.0.0
UID0:		PWD0:	••••
UID1:	0	PWD1:	•
GkOrProxy:	sipauth.deltathree.com	UseLoginID:	1
LoginID0:		LoginID1:	0
AltGk:	natrelay.deltathree.com	AltGkTimeOut:	0
SIPRegInterval:	300	MaxRedirect:	10
SIPRegOn:	1	NATIP:	0.0.0.0
SIPPort:	5060	MediaPort:	16384
OutBoundProxy:	0	NatServer:	stun1.deltathree.com
NatTimer:	0x0000001d	MsgRetryLimits:	0x00000000
SessionTimer:	0x00000000	SessionInterval:	1800
MinSessionInterval:	1800	DisplayName0:	Line 0
DisplayName1:	Line 1	LBRCodec:	6
AudioMode:	0x00200020	RxCodec:	6
TxCodec:	6	NumTxFrames:	2
CallFeatures:	0xffffffff	PaidFeatures:	0xffffffff
CallerIdMethod:	0x0001a8a0	FeatureTimer:	0x00000000
FeatureTimer2:	0x0000001e	Polarity:	0x00000000
ConnectMode:	0x17190484	TimeZone:	20
NTPIP:	192.5.41.40	AltNTPIP:	198.123.30.132
DNS1IP:	213.137.73.254	DNS2IP:	208.170.171.93
TOS:	0x000000b8	SigTimer:	0x03c18364

Broadband Phone
$10.95
iConnectHere
www.iconnecthere.com
CPU CPU C

CPU Ranking: 0 = Absolutely Worthless CPU CPU C 2.5 = Absolutely Average CPU CPU CPU CPU CPU 5 = Absolutely Perfect

them to the network), iConnectHere made me configure the gateway manually, connecting to it with a Web browser to adjust parameters. It was fun for a geek like me, but I wouldn't expect a novice to do it, and it certainly diminishes the easy-to-use-as-a-telephone advantage.

The quality was very good for stateside calls, with no noticeable latency. Some people I spoke with thought the call could be from a traditional phone. On my end, however, the calls sounded like a digital mobile phone. In addition, calls to Germany wouldn't go through, returning a "Your call to this destination could not be completed" message or dead air. When searching for help at the company's Web site, I only received returned server errors.

iConnectHere includes a handful of calling features, although fewer than some other services. Features include voicemail, caller ID, and the ability to forward incoming calls to any other phone number. The company offers several calling plans, including a per-minute plan with no flat monthly charge. Under that plan, calls to the continental United States are 2.3 cents per minute. For $10.95 per month, you get 1,000 minutes to North American destinations. This was the only service that didn't offer an unlimited use plan.

Net2Phone VoiceLine

Net2Phone's VoiceLine service isn't the cheapest, but it was reliable and feature-filled. Call quality for stateside and international calls was generally quite good, clear, and loud. During one call, I could hear my own voice echo, which was distracting but didn't reoccur when I called again. VoiceLine also transmitted my caller ID information to Germany.

Calling features include voicemail, which you can retrieve on the phone or the Web; call waiting; caller ID; call return (also known as *69); three-way calling; and call blocking, which lets your line automatically reject anonymous callers (or calls from your ex). You can manage these features via the Web or use the phone's touch-tone keys.

At $34.99 per month for unlimited calling to the United States and Canada, VoiceLine is tied as the most expensive service here. For $14.99 a month, you can get 500 minutes of outgoing calls or 300 minutes for $9.99. Local phone numbers are only available in 11 area codes in five states, which is sparse compared to some services, but U.S. and UK toll-free numbers are also available for an extra fee.

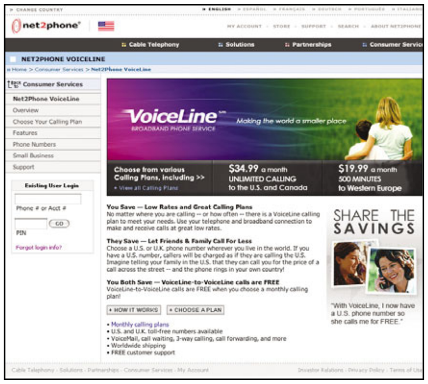

Choose from various Calling Plans, including >>
View all Calling Plans
$34.99 a month UNLIMITED CALLING to the U.S. and Canada
$19.99 a month 500 MINUTES to Western Europe

You Save — Low Rates and Great Calling Plans
No matter where you are calling -- or how often -- there is a VoiceLine calling plan to meet your needs. Use your telephone and broadband connection to make and receive calls at great low rates.

They Save — Let Friends & Family Call For Less
Choose a U.S. or U.K. phone number wherever you live in the world. If you have a U.S. number, callers will be charged as if they are calling the U.S. Imagine telling your family in the U.S. that they can call you for the price of a call across the street -- and the phone rings in your own country!

You Both Save — VoiceLine-to-VoiceLine calls are FREE
VoiceLine-to-VoiceLine calls are FREE when you choose a monthly calling plan!

- Monthly calling plans
- U.S. and U.K. toll-free numbers available
- VoiceMail, call waiting, 3-way calling, call forwarding, and more
- Worldwide shipping
- FREE customer support

"With VoiceLine, I now have a U.S. phone number so she calls me for FREE."

Cable Telephony · Solutions · Partnerships · Consumer Services · My Account · Investor Relations · Privacy Policy · Terms of Use

VoiceLine
$34.99
Net2Phone
www.net2phone.com

CPU CPU CPU CPU

Packet8

Packet8 was one of the least expensive services I tested to offer unlimited calling. For $19.95 a month, you get unlimited calling in the United States and Canada. For $79.90 a month you can get unlimited calls to Europe and Asia. Features include voicemail, call waiting, call forwarding, three-way conferencing, and caller ID. This is a small feature set compared to other services, but it's probably enough for most people. There are virtually no features on the Web site, however.

The quality of U.S. calls was fair, akin to digital cellular service. One call had latency, which led to awkward pauses in the conversation and participants talking over each other. On another call I could hear fine but the other participant asked me to repeat myself several times. On yet another call, dropouts meant I missed a few words here and there. I also couldn't complete my call to Germany, only getting a fast busy signal instead.

Although its features are sparse and voice quality is far from perfect, Packet8's low monthly fee could make it a good choice for chatterboxes on a budget.

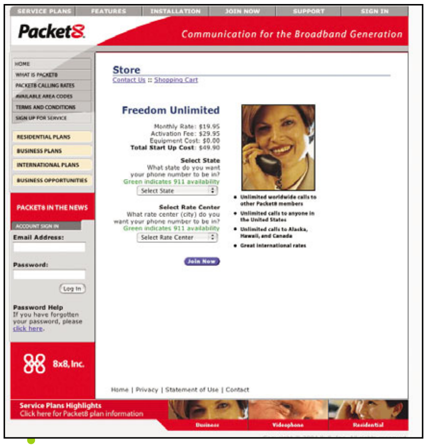

Store
Contact Us :: Shopping Cart

Freedom Unlimited

Monthly Rate: $19.95
Activation Fee: $29.95
Equipment Cost: $0.00
Total Start Up Cost: $49.90

Select State
What state do you want your phone number to be in?
Green indicates 911 availability

Select Rate Center
What rate center (city) do you want your phone number to be in?
Green indicates 911 availability

- Unlimited worldwide calls to other Packet8 members
- Unlimited calls to anyone in the United States
- Unlimited calls to Alaska, Hawaii, and Canada
- Great international rates

Email Address:
Password:

Password Help
If you have forgotten your password, please click here.

Home | Privacy | Statement of Use | Contact

Service Plans Highlights
Click here for Packet8 plan information

Packet8
$19.95
Packet8
www.packet8.net

CPU CPU CP

voiceglo

Compared to the other services, voiceglo was disappointing, primarily because calls were unexpectedly cut short. Call quality to U.S. destinations was fair. Voices tended to sound fuzzy, and one person I spoke with noticed the start of each sentence being slightly cut off. An unpleasant surprise came 51 minutes into one conversation when the call inexplicably ended with a busy signal. A later conversation was also cut off after 10 minutes. My five-minute call to Germany sounded great, however, with low latency.

voiceglo offers two calling plans for residential users. A $19.99 Regional plan

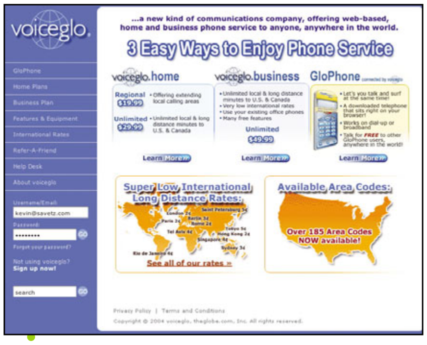

...a new kind of communications company, offering web-based, home and business phone service to anyone, anywhere in the world.

Regional $19.99 · Offering extending local calling areas
Unlimited $29.99 · Unlimited local & long distance minutes to U.S. & Canada

- Unlimited local & long distance minutes to U.S. & Canada
- Very low international rates
- Use your existing office phones
- Many free features

Unlimited $49.99

- Let's you talk and surf at the same time!
- A downloaded telephone that sits right on your browser!
- Works on dial-up or broadband
- Talk for FREE to other GloPhone users, anywhere in the world!

Username/Email: kevin@savetz.com
Password: ••••••••
Forgot your password?
Not using voiceglo? Sign up now!
search

Privacy Policy | Terms and Conditions
Copyright © 2004 voiceglo, theglobe.com, Inc. All rights reserved.

voiceglo
$19.99 (Regional)
voiceglo
www.voiceglo.com

CPU CPU

CPU Ranking: 0 = Absolutely Worthless CPU CPU C 2.5 = Absolutely Average CPU CPU CPU CPU CPU 5 = Absolutely Perfect

Dial In On A Service

The seven broadband phone services tested in this roundup have numerous similarities, but also many differences. Here is how the services compare.

Service	Pros	Cons	Cost/month for unlimited U.S. and Canadian calls	URL	CPUs
AT&T CallVantage	Excellent voice quality for U.S. calls	Poor quality for international calls; difficult setup; most expensive	$34.99 (U.S. only; Canada adds 5 cents per minute)	www.usa.att.com/callvantage	CPU CPU CPU
Broadvox Direct	Good voice quality	Incoming call delivery unreliable	$29.95	www.broadvoxdirect.com	CPU CPU
iConnectHere Broadband Phone	Good voice quality	Manual configuration	No unlimited use plan; $39.95 for 3,000 minutes	www.iconnecthere.com	CPU CPU C
Net2Phone VoiceLine	Very good voice quality	Numbers available in only 11 area codes	$34.99	www.net2phone.com	CPU CPU CPU CPU
Packet8	Least expensive service	Fair voice quality; few extras	$19.95	www.packet8.net	CPU CPU C
voiceglo	Nothing special	Calls cut off without warning; fair voice quality	$29.95	www.voiceglo.com	CPU CPU
Vonage	Top-notch voice quality; many features	No major flaws	$29.99	www.vonage.com	CPU CPU CPU CPU C

provides unlimited local minutes and long distance U.S. and Canadian calls for 3.9 cents per minute. I don't see any benefit here, though, as plenty of traditional long distance carriers charge less. (See www.telcompare.com for a list.) voiceglo's $29.99 Unlimited plan with unlimited U.S. and Canadian minutes makes more sense.

voiceglo's features are good but don't rock the house with extras. Included is voicemail, caller ID, call waiting, call forwarding, speed dial, and reviewing call logs online. When I signed up, I didn't receive a password to check logs or do online account maintenance. I used the site's Forgot Your Password? function, but when it arrived 24 hours later, it didn't work. I never could log in to manage my account.

Vonage

Partly due to aggressive marketing, Vonage is perhaps the company most associated with broadband phone calling. Its reputation is well-deserved, as it offers excellent call quality and many features for a fair price.

The audio quality of domestic and overseas calls was consistently top-notch, without any lag or digital artifacting. No one I spoke with guessed that I wasn't using a regular phone. Caller ID information did not make it to my German friend, but that's a minor issue.

Vonage offers several plans for residential and business users. The $29.99 per month Premium Unlimited plan includes unlimited U.S. and Canadian calls. The long list of features includes voicemail, call waiting, call forwarding, repeat dialing, caller ID, three-way calling, and call transfer. You can also get a new phone number in the city of your choice or transfer your existing POTS number to Vonage.

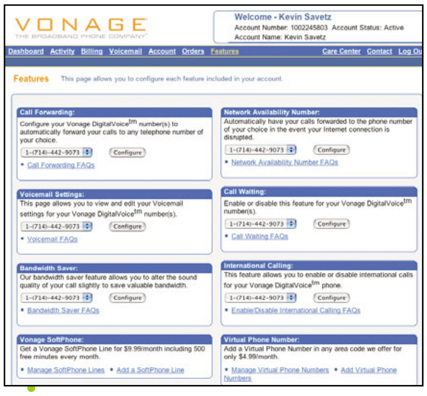

Welcome - Kevin Savetz

Account Number: 1002245803 Account Status: Active

Account Name: Kevin Savetz

Features This page allows you to configure each feature included in your account.

Call Forwarding: Configure your Vonage DigitalVoice^{tm} number(s) to automatically forward your calls to any telephone number of your choice.
1-(714)-442-9073
- Call Forwarding FAQs

Network Availability Number: Automatically have your calls forwarded to the phone number of your choice in the event your Internet connection is disrupted.
1-(714)-442-9073
- Network Availability Number FAQs

Voicemail Settings: This page allows you to view and edit your Voicemail settings for your Vonage DigitalVoice^{tm} number(s).
1-(714)-442-9073
- Voicemail FAQs

Call Waiting: Enable or disable this feature for your Vonage DigitalVoice^{tm} number(s).
1-(714)-442-9073
- Call Waiting FAQs

Bandwidth Saver: Our bandwidth saver feature allows you to alter the sound quality of your call slightly to save valuable bandwidth.
1-(714)-442-9073
- Bandwidth Saver FAQs

International Calling: This feature allows you to enable or disable international calls for your Vonage DigitalVoice^{tm} phone.
1-(714)-442-9073
- Enable/Disable International Calling FAQs

Vonage SoftPhone: Get a Vonage SoftPhone Line for $9.99/month including 500 free minutes every month.
- Manage SoftPhone Lines
- Add a SoftPhone Line

Virtual Phone Number: Add a Virtual Phone Number in any area code we offer for only $4.99/month.
- Manage Virtual Phone Numbers
- Add Virtual Phone Numbers

Vonage
$29.99
Vonage
www.vonage.com
CPU CPU CPU CPU C

Even better, Vonage offers virtual phone numbers, which let you have phone numbers in more than one city, all of which ring your phone. So, you could have a main phone number in California and a virtual phone number in New York, letting friends and associates in both cities call you with a local call. Each virtual number costs $4.99 per month. There's also an unusual Bandwidth Saver feature that lets you use less bandwidth (as little as 30Kbps) in exchange for reduced voice quality.

It's For You

A broadband phone service can be quite useful as a second home line, business line, or even as a replacement for your current phone service. At $20 to $35 per month for unlimited U.S. calls, these services can save you money if you make a lot of long distance calls. The services do charge extra for international calls, although typically less than traditional long-distance carriers. When you factor in the basic cost of a phone line and any extras you pay for, broadband phone services can be a bargain. ▲

by Kevin Savetz

CPU Ranking: 0 = Absolutely Worthless CPU CPU C 2.5 = Absolutely Average CPU CPU CPU CPU CPU 5 = Absolutely Perfect

Immaculate Wi-Fi Reception

The 802.11g USB 2.0 adapter with twist-and-shout flexibility!

IOGEAR's ultra fast Wireless-G 802.11g USB 2.0 adapter can help you connect to a wireless network with "higher power." It bends and twists to ensure optimal reception so you can work at up to 54Mbps at home or anywhere a hotspot is available.

Installation takes just moments. All IOGEAR Wireless-G products offer the latest security and are backward compatible with 802.11b.

So "retire the wire" without sacrificing Internet speeds - pick up IOGEAR's line of Wireless-G products today!

GWU513

Wireless-G to USB 2.0 Flex Adapter

- Uses USB port, not cardbus or PCI slot
- No configuration hassles
- Flexes and pivots for optimal reception

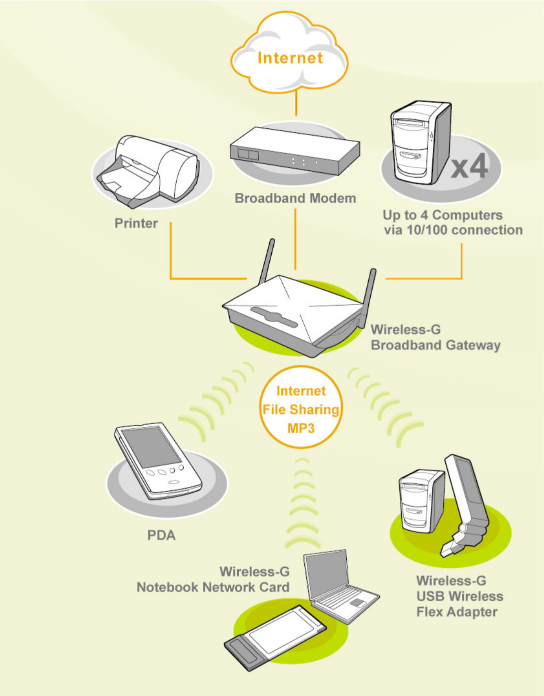
Internet
Broadband Modem
Printer
x4
Up to 4 Computers via 10/100 connection
Wireless-G Broadband Gateway
Internet File Sharing MP3
PDA
Wireless-G Notebook Network Card
Wireless-G USB Wireless Flex Adapter

GWA501

Wireless-G Broadband Gateway/Print Server

- Fast setup in just a few simple steps
- Includes everything for effortless installation
- Built-in firewall, supports WPA and WEP
- Built-in print server shares your printer with all connected computers

GWP511

Wireless-G Notebook Network Card

- Low-profile high-gain antenna achieves favorable sleek design
- Included Software Access Point allows you to turn your notebook into a wireless Access Point

Retire the Wire.

IOGEAR Inc.
23 Hubble Irvine, CA 92618 P: 949.453.8782 ext. 2011 http://smc.iogear.com

Tech Data

IOGEAR®
New Thinking, New Style

© 2004 IOGEAR, Inc. All rights reserved. All other trademarks are the property of their respective owners.

by Kevin Savetz

MightSoft Audio Editor Pro 1.40

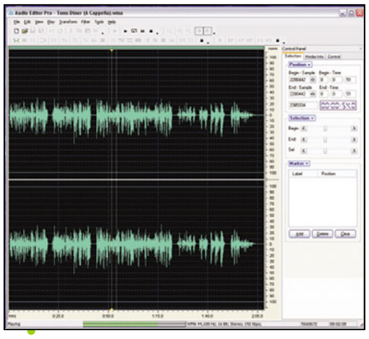

Audio Editor Pro 1.40
- $39.95
- MightSoft
- www.mightsoft.com

CPU CPU CPU C

A good sound editor can be a truly useful and fun addition to your software toolbox. You can use it to trim the annoying introduction from an otherwise-great MP3 song, create your own Windows error sounds, or splice together a unique answering machine message. For those tasks and countless others, Audio Editor Pro does an admirable job.

The heart of the app is its waveform editor, which presents a visual representation of your sound file. If you've used any sort of waveform editor before, you'll be able to jump headlong into AEP. By selecting a portion of the audio, you can manipulate it in various ways. The Transform menu includes various tweaks, including Reverb, Fade, Flanger, and Stretch. The Filters menu provides Low and High Pass, Low and High Shelf, Ban Pass, and Notch filters.

A Spectral view lets you see the tonal range of your samples. When capturing samples from a mic or line-input source, you can have the app automatically apply filters to remove tape hiss or other spurious noise. I do have a few quibbles, however. You can't perform most functions, even something as simple as zooming in on a selection, while the program is playing back the audio. Also, expanding a selection to an area outside the current view could be easier, as you can't click and drag to scroll the view.

AEP works with MP3, WAV, WMA, Ogg Vorbis, plus other audio formats. Although it can open and save those formats directly, the package includes a simple app for converting between audio file formats, including batch conversions. There's also a basic CD ripper, but it didn't work well, failing to recognize a FireWire optical drive. The CDDB feature also didn't work for me.

Despite a few hiccups, AEP is a worthy sound editor. It doesn't demand a fast PC. A 400MHz CPU will get the job done, and you can "sample" the program for 15 days before buying it. ▲

TuneUp Software TuneUp Utilities 2004

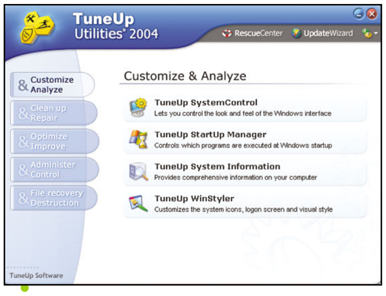

TuneUp SystemControl
Lets you control the look and feel of the Windows interface

TuneUp StartUp Manager
Controls which programs are executed at Windows startup

TuneUp System Information
Provides comprehensive information on your computer

TuneUp WinStyler
Customizes the system icons, logon screen and visual style

TuneUp Utilities 2004
- $39.99 ($24.99 upgrade)
- TuneUp Software
- www.tune-up.com

CPU CPU CPU

Keeping Microsoft Windows happily running isn't fun for anyone. We may be power users, but we aren't masochists. Dealing with Registry errors and other maintenance issues are chores no one likes. TuneUp Utilities 2004 can take some of the pain away, plus it provides some cool Windows customizations.

Some of the program's features parallel Norton SystemWorks'. For example, TuneUp will search for invalid Registry entries, search for missing files your applications need, and remove space-eating temporary files. A 1-Click Maintenance tool will quickly run through the tune-up regimen. The tool indicates how many problems it finds, but it doesn't provide specifics about what those issues are. You'll have to run each utility by hand for that information.

Several features mirror functions built into Windows. For example, when memory gets tight, the memory optimizer frees RAM by moving unneeded programs to the paging file. It's difficult to say if it does a better job than Windows' default paging system, however. There's also a Registry editor and process manager, both similar to those included in Windows, although with some additional functions.

The app also includes an Uninstall Manager that removes apps you no longer need; a DiskCleaner that removes backup, log, Internet cache, and temporary files; an Undelete utility that restores files; and Shredder, a DoD-compliant tool for securely deleting files, making sure other users can't recover your deleted secrets.

A WinStyler feature provides clever functions for customizing the Windows interface. You can use it to modify Windows' visual style; change the image on the Windows login screen; and replace the default My Computer, Recycle Bin, and other icons with graphics of your choice.

TuneUp Utilities was stable and speedy in my tests, repairing Registry problems more quickly than Norton SystemWorks. It is lacking as a complete system tune-up solution, however, failing to include file and disk-repair functions. In addition, it doesn't make a strong case for why some of its tools are better than Windows' built-in functions. ▲

CPU Ranking: 0 = Absolutely Worthless CPU CPU C 2.5 = Absolutely Average CPU CPU CPU CPU CPU 5 = Absolutely Perfect

by Chris Pirillo

Remotely Interested?

Chris Pirillo is addicted to rhymes, written in Helvetica or scribbled down in Times. Lockergnome.com is the way he makes his dimes. He's been known to coddle lemons but much prefers the limes. Of all the numbers known to us, his favorites are the primes. A suburb of Des Moines is the city known as Grimes. When he's forced to disconnect, his frustration level climbs. Of all the things he's ever done, none were legal crimes. He's also never ever met the girl named LeeAnn Rimes.

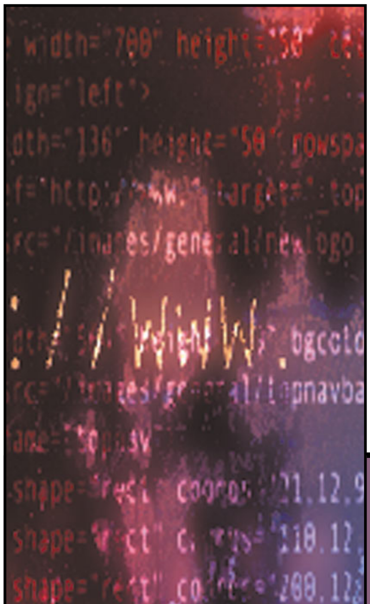

Phone rings. Answer. Computer question comes. Solutions galore return. Same computer question remains. Lather, rinse, and repeat until hair falls out.

When it comes to dealing with computer issues from afar, seeing is believing. No matter how much you trust dear old dad, buying what he says happened is a leap of technical faith. The best thing to do for relatives, friends, and strangers is introduce them to remote control software.

The choices are as varied as the situations that necessitate them. A good starting place is VNC (www.realvnc.com), the open-source system UK-based RealVNC now maintains. Some would argue that UltraVNC is better (www.ultravnc.com). Regardless, there's a VNC client for just about any system, including Novell Netware servers, HP-UX, OS/2, QNX, Palm, and Pocket PC. VNC is the king of cross-platform, largely due to its open-source nature, allowing many developers across the world to port, extend, and hack to suit their particular needs. There's even a port for the Apple Newton!

Spending hundreds to bail out relatives may indeed be overboard . . .

If you prefer a more refined process, Symantec has a leading commercial app in pcAnywhere. Thankfully, it's far more organized than Symantec's example of atrocious Web site navigation. Assuming you find what you're looking for, pcAnywhere performs well, has many administration goodies, and sports the most flexible authentication. However, you'll pay an equally daunting price ($199.95) for what amounts to a Windows-centric focus.

If you need to mix Windows and Mac OS from here to Timbuktu, Netopia (www.netopia.com/software/products/tb2) has just the thing. The usual remote control and file-transfer operations make up the bulk of Timbuktu, but it also includes a voice intercom gadget. Save those long-distance charges and have the parents explain their misstep through the Internet. It's not quite a landline, but it's acceptable.

The other notable commercial package is CrossTec's NetOp (www.crossteccorp.com/netopremote). It takes many of pcAnywhere's best traits—centralized administration, multiple authentication methods, and great performance—and adds fantastic cross-platform support, including Windows, Mac, Linux, Solaris, OS/2, and Pocket PC. You can also use an ActiveX component via IE to control systems running the full host software app, expanding the fluid nature of remote administration. Of all commercial titles, NetOp is the most impressive in scope but is just less than ($189) pcAnywhere.

Spending hundreds to bail out relatives may indeed be overboard, but something like VNC still requires a separate installation and follow-up attention with regard to security and any changes in network configuration. Using an online service can alleviate that burden. GoToMyPC (www.gotomypc.com) is the chief player here. It's most attractive to those needing access to various systems over the Internet and when a dedicated software package is too restrictive in licensing terms. GoToMyPC has personal/pro/corporate licensing plans, putting the ball in your court. It sounds expensive at $19.95 a month for the single-user account, but you can control multiple systems. Commercial software requires an explicit license for every PC you want to access. If you're the only geek in the family, I don't have to do the math for you.

Microsoft added Terminal Services functionality to WinXP, and no discussion of distance assistance would be complete without some coverage. There are two distinct courses you can follow here. By default, WinXP enables Remote Assistance, which depends on the user to initiate a request for aid and requires steps many users won't understand or want to deal with. The hoops dictate creating a session "invitation" and sending it via Windows Messenger (requiring a Passport ID, naturally) or email to the desired party who must use Outlook or Outlook Express. If the settings aren't specified correctly for the network environment, it doesn't work.

For WinXP Pro users, the alternative is enabling Remote Desktop. This easier method was left out of WinXP Home ironically, and it still has a major downside. Once you take control, the user sees only a box that informs them the workstation has been locked. If they unlock it, your session is cut instantly. You can't work together on a problem, so coordination of the asynchronous process is necessary. Frankly, it's easier to use VNC if free is a must. ■

You can dialogue with Chris at chris@cpumag.com.

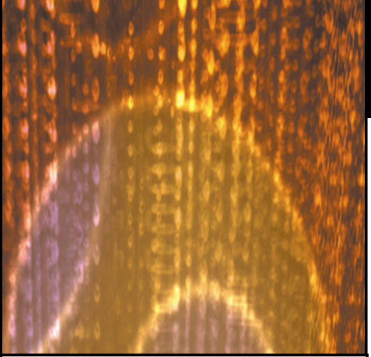

Open Sauce

by Pete Loshin

Going Open Source Is Cheaper

Pete Loshin, former technical editor of software reviews for Byte Magazine (print version), consults and writes about computing and the Internet. He also runs www.linuxcookbook.com. He owns shares of both Microsoft and Red Hat and believes that Windows isn't for everyone, but neither is Linux.

Ever wonder what it costs to equip a Windows PC with commercial software for business use? About $1,500.

Upgrades from older versions of some commercial apps can save a few hundred. Big corporations buy in bulk for huge savings. Savvy shoppers can undoubtedly find cheaper sellers than Amazon.com. But I'd rather shop the hardware and walk away with enough to buy the components for two or three actual PCs running Linux than pay for all that software.

Big consulting firms report drastically different results about the cost of buying and using proprietary software, but big commercial software publishers pay for those reports.

My first task was to figure out what's mandatory for a Windows PC. Not having used Windows in years, I asked some experts: editors of several computing magazines and subscribers to a secret (oops) mailing list for computer and technology experts. In other words, people who make their livings as experts about computing. Here's what they told me about must-have software, along with pricing.

[* indicates Amazon.com's price, often including rebates or other offers, as of June 4, 2004. ** indicates a publisher's direct price. One note: I know of no free tax-prep software. If you have to use it, you'll have to borrow someone else's Windows PC. Sorry.]

Operating system: WinXP Pro ($264.99* for the full product; the upgrade is $175.99*). The nonupgrade Home Edition is $176.99* (upgrade $89.99*), but XP Pro touts "faster startup, logon, and launch times." Does that mean the Home Edition is extra slow? Why is there a difference? Anyway, Pro's networking and security features are more like standard Linux.

> My first task was to figure out what's mandatory for a Windows PC.

Office suite: Microsoft Office Pro ($409.99*) supports industry-standard XML, not just Microsoft-proprietary XML, and includes Word, Outlook, Excel, Access, PowerPoint, and Publisher.

Microsoft FrontPage 2003 for Web development ($159.99*) and Adobe Acrobat 6.0 Standard Edition for creating PDF files ($249.99*) both go in the cart. You may get Adobe's Photoshop Elements in a bundle with a camera or printer, but it, bundled with Photoshop Album, seemed cheap at $77.99* (with $30 rebate). Also noted was Jasc Paint Shop Pro ($52.99* with $30 rebate) for graphics and photo editing.

Personal finance and tax software were two product categories high on everyone's lists. I picked Intuit's Quicken 2004 Deluxe for $49.99*. Microsoft's Money lists at $59.99 but is discounted $15, and coupled with a $20 rebate (which gets you on Microsoft's junk mail list) the price is $24.99*. It can't be much good if it's so cheap, right?

If you're using Windows, you have to have some security software. Norton Internet Security 2004 Professional (for "AntiVirus, Firewall, AntiSpam, Spyware, Privacy, Productivity Control") at $51.99* after a $30 rebate is one popular option. Or you can get Norton Internet Security 2004 ("Parental Control" instead of "Productivity Control") for $9.99* (after a $50 rebate).

One respondent advocated a pop-up killer/ad stopper/spyware blocker like AdsGone ($19.95**; adsgone.com), while others pointed to utilities like ExplorerPlus from Novatix ($39.95**; novatix.com) and ZipMagic from Alladin Systems ($39.99**; aladdinsys.com). RecordNow! Deluxe Suite ($69.99**; sonic.com) was highly recommended for ripping/burning CD/DVDs.

Microsoft Flight Simulator 2004 is one of Amazon's best sellers at $24.99*, and many respondents said gaming is important to them. I didn't include sound/video-editing software because I couldn't get a good consensus on what was absolutely necessary.

Here's the bottom line: $1,512.79 for one PC. Or enough to buy the parts to build one great, two average, or three budget PCs. Maybe I have the list wrong; maybe Amazon.com charges too much. Email me at must.have.software@loshin.com, and tell me what to drop and add. In two months I'll have a final list, plus open-source alternatives to help keep your money in your own pocket. ■

Get saucy with Pete at pete@cpumag.com.

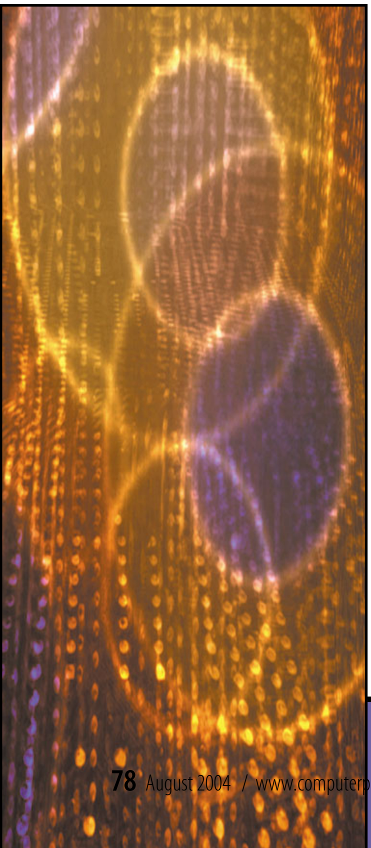

Hearing is Believing

PC WORLD
WORLD CLASS
2003

revolution 7.1

High-Definition Surround Sound for Your PC

Revolution 7.1

- PCI solution for high-definition surround up to 7.1
- 24-bit/192kHz output up to 8 channels
- Professional 24-bit/96kHz recording
- Real-time stereo-to-surround conversion
- Incredible 107dB SNR provides pristine audio

Sonica Theater

- USB solution for surround up to 7.1
- Up to 24-bit/96kHz audio output
- Professional 24-bit/96kHz recording
- Real-time stereo-to-surround conversion

TRUE 24 bit 192 kHz High Definition

M-Audio's audio cards deliver immersive surround sound on your computer, rivaling that of the best movie theaters. Games, movies, and music—even your stereo CDs and MP3s—all sound more alive and intense with M-Audio surround solutions. Now you can enjoy amazing new cutting-edge audio standards like 7.1 surround sound and high-definition 24-bit/192kHz fidelity. Support for popular surround formats including Dolby Digital and DTS, as well as gaming standards like Sensaura, EAX, DirectSound and A3D, insures that you'll hear everything exactly the way it was intended.

Only one name is found in more studios and is used by more music professionals than any other: M-Audio. Now we've distilled all that experience into the most powerful, high-fidelity PC surround sound solutions available. In fact, PC World recently gave the Revolution 7.1 its highest accolade–the coveted World Class Award 2003.

Get M-Audio. Get M-Powered.

available at
COMPUSA®
WHERE AMERICA BUYS TECHNOLOGY.℠

M-AUDIO
www.m-audio.com

Not Just A FLASH In The Pan

Macromedia's Design Tool Hits The Big Time

Most of us still think of Macromedia Flash as a tool for making those annoying animated Web site intros or intrusive pop-up ads that we quickly skip through to get to the *real* home page. Not so fast, Internet snob. Over the years, this once-simple design tool for making online rich-media presentations, Web sites, and Web cartoons has grown into a very powerful design and programming platform that has definitely burst out of the site intro.

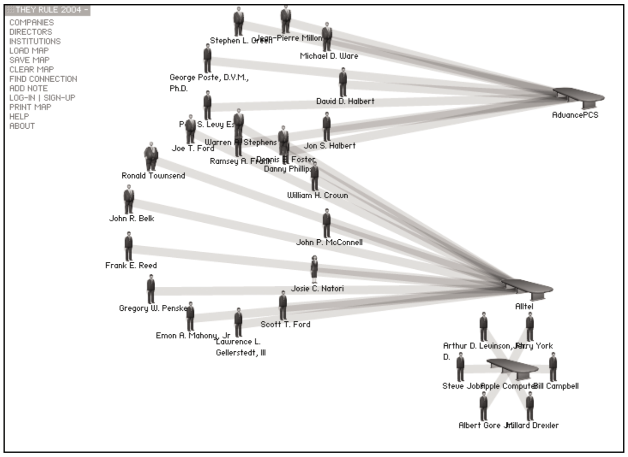

Artful Flash programmers are using the graphics-oriented tools to create new ways of viewing information and making connections. At TheyRule.net, you can visualize boards of directors at the world's top companies to see their common members.

In 2004, the chic and renowned Sundance Film Festival used Flash to stream its independent films online (www.sundanceonlinefilmfestival.org) because one of the platform's new features is the ability to run streaming video within a Web page. "One challenge we wanted to overcome was drawing people into a more artistic way of viewing films on the Internet," says Amy Ozski, executive producer, Sundance Online Film Festival. "Flash allowed for the incorporation of moving images [into Web pages], not just still images."

At TheyRule.net, designers are using Flash to visualize information in novel ways, such as how many major U.S. companies share a relatively small pool of board members. Brooke Burgess, creator of the 24-episode "cinematic literature" opus entitled "Broken Saints" (www.brokensaints.com), says he "uses Macromedia's Flash to tell a creepy and thought-provoking story." This atmospheric, spiritual tale is about to be transferred to DVD for general release in the fall. Much like film itself, Flash has become a versatile medium that is being put to use both for creative expression and business productivity.

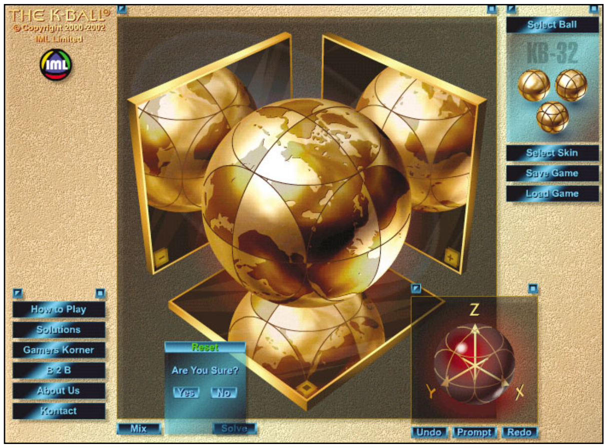

Flash is being used to make games that look and feel like art. K-Ball takes the Rubik's Cube to a new level with mirrors, a skinnable spherical puzzle, and something we always wanted in our Rubik's cube: a button to restore the puzzle to its starting point.

Some purist defenders of Web usability like to denounce Flash as a main cause for Webmasters overdesigning their sites with unnecessary bells and whistles and cute but obtuse navigation schemes. Nevertheless, most designers are drawn to Flash because it can go so far beyond HTML. The Flash Player plug-in required to run Flash's SWF file format is the closest thing to a universal platform for online rich media, second only to HTML code itself in its compatibility with desktops.

With a reach of 481 million users, "ninety-eight percent of Internet-connected desktops have the Flash Player," says Jeff Whatcott, vice president of product marketing at Macromedia. Over the years, Flash has gained audio capabilities and the abilities to hook into back-end databases and now to stream both video and MP3s all within a Web page and do so in player software "that's still only 477KB," says Whatcott. And because it is a vector-based tool, it adjusts well to different browsers and display resolutions so that the developer can be assured the user is seeing the site or presentation as it was made.

The Flash Revolution

In part because Flash is so easy to learn for nonprogrammers, it has generated an amazing range of creative products. The semiannual FlashForward Conference (www.flashforward2004.com) is a three-day gathering of Flash fans at all levels of expertise that culminates in a series of awards for the best programming across 15 categories. FlashForward receives thousands of submissions, from site intros to games and business applications to experimental Flash-based art. Sites such as Bit-101 (www.bit-101.com) use the tool to make a Web site interface that looks and functions exactly like a customizable computer desktop whose elements pop up as movable, resizable windows. Unique audio-visual presentations such as Luis Escorial's "The Beauty of Details" at his One Thousand Dreams site (www.1000dreams.com) is an animated interactive poem where mouse clicks move the artistic expression in different directions. Because of the way it blends multimedia seamlessly and allows users to interact with on-screen vector-based images, Flash has become *the* platform for digital media experiments.

Flash has long been used to make Web-based games, but some developers, such as Cartoon Network (www.cartoonnetwork.com), are bringing it to the next level. Its gToons exploited Flash's XML handling, as well as compatibility with Java Servlet environments, to build a sophisticated head-to-head trading card game. Visitors could both play and communicate with each other on the site in real-time, which brought in 400,000 sign-ups for 5.3 million plays in two months, according to the company.

One of the most intriguing combinations of Flash technology with creative energy comes from Boston musician, composer, and programmer Alexander Chen, whose Sonata For The Unaware (www.carbonatedjazz.com) generates musical scores from the videos of Philadelphia commuters trudging to and from work. He used Macromedia Director to process the video. "This code automatically analyzes the videos pixel by pixel and outputs raw data, which is read by Flash. (See "Beyond The Intro" for more.) Flash then generates the music from the raw data while the videos play," says Chen.

While games and episodic Web animation are among the best-known creative uses of Flash, Kymberlee Weil, organizer of FlashForward, says that the platform has now become an industrial-strength tool that is driving businesses. "It used to be about Web sites, and now it is about applications," she says. Flash has become a favorite tool of museums and online training because it is so visually oriented and interactive. A resource for the visually impaired, Lighthouse International (www.visionconnection.com) designed online training for visual rehabilitation assistants, which usually required them to come to headquarters in New York City. Patrick Benson, CIO, Lighthouse International, was able to use text-to-speech tools, animated illustrations of rehab techniques, and interfaces that accommodated visually impaired users. "The economics of the distance learning course we created using Flash are amazing," Benson says in a Flash case study paper. "The savings of just airfare and temporary accommodations . . . are tremendous."

And then some Flash executions (perhaps the most involving ones) are at the bleeding edge creatively and conceptually. Gicheol Lee's Typorganism (www.typorganism.com), where you play interactively with letters you type on the screen or with headlines and images pulled down in real-time from CNN.com, is a good example. Lee calls these "communications experiments" based on the notion that "type is a lifeform." Right.

by Steve Smith

Up In A Flash

Flash is among the easiest and cheapest development tools a novice can learn. "I don't have a programming background," says Kymberlee Weil, who has herself authored two books on learning Flash. For starters, anyone can download a 30-day free demo of the $499 Flash MX 2004 from www.macromedia.com. Also helpful for the beginner is the site's On Demand section of video tutorials that bring the novice up to speed on the basics of the interface. And Macromedia hosts a robust Flash Developer Center (www.macromedia.com/devnet/mx/flash) with detailed tutorials for all aspects of making animations, incorporating video, and connecting Flash front ends to existing databases.

From there, go online or get a book. Sites such as Lynda.com offer online video tutorials for Flash beginners. LearnFlash.com has downloadable tutorials for specific tasks such as creating animations or site intros in Flash. The number of books teaching Flash are too numerous to mention, but among the best we have seen are "Macromedia Flash MX 2004 Hands-On Training" by Rosanna Yeung ($44.99; Peachpit Press), "Macromedia Flash MX for Windows and Macintosh" by Katherine Ulrich ($24.99; Peachpit Press), and "SAMS Teach Yourself Macromedia Flash MX 2004 in 24 Hours" by Phillip Kerman ($29.99; SAMS).

And if you are looking for a creative jump-start, "Come to the FlashForward Conference," says organizer Weil. The three-day semiannual affair is the Cannes Film Festival of Flash. It alternates between New York and San Francisco with seminars and workshops for all levels of expertise. FlashForward culminates in an awards ceremony where Flash is treated much like film, with trailers showing nominated work on the big screen before winners are announced. "Just being around these people and seeing the work they are doing is so inspiring," she says. Or, take a look at our own gallery of cutting-edge Flash work in the "Beyond The Intro" section of this article to see how the former tool of the annoying animated site intro has now become a full-fledged platform for creativity and productivity. ▲

Infinite Loop

DDR, Weight Loss & You

Sorry, folks. We're not talking about double-data-rate memory. But if you're a frequent visitor to the arcade at your local mall, you probably know what we're talking about. Techno music, psychedelic colors, feet a-flyin' like Flatley—it's the Dance Dance Revolution revolution. In addition to the strong arcade presence, Konami's popular series has already sold 1 million home copies in the United States and 6.5 million worldwide. The game's appeal isn't limited to the average gamer or music and dance aficionados, however. A new group has taken interest in the game: the fitness buffs. A recent AP article interviewed individuals who had lost anywhere from 60 to 150 pounds as a result of habitual game play and better dieting. ▲

Source: Associated Press

Beyond The Intro

Flash programming has gone way beyond animated intros, annoying pop-up ads, and relentless Tetris clones. Here is just a sampling of cutting-edge uses of the Macromedia tool in the fields of art, education, gaming, and business applications. But don't miss out on the fun of discovering the creative energies Flash has unleashed onto the Web. Some of the best examples are at the FlashForward Film Festival (www.flashforward2004.com) and Macromedia's own Showcase (www.macromedia.com/cfusion/showcase/index.cfm).

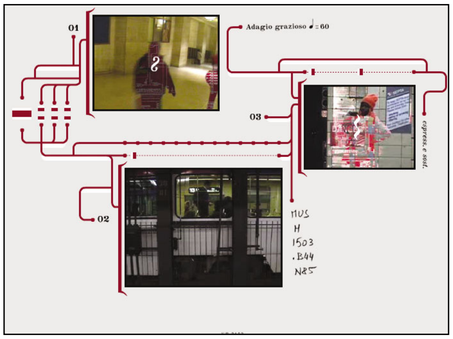

Sonata For The Unaware • www.sonatafortheunaware.com

Using three video windows of Philadelphia commuters blithely walking past security cameras, musician/programmer Alexander Chen tracks their physical movements and translates them into notes in a sonata. Chen describes it as "human bodies becoming musical markings on the pages of a cranking machine." He used multiple musical engines he had designed in Flash, including a moog synthesizer and a drum sequencer. The piece also demonstrates one of the most popular uses of Flash's newer capabilities to blend streaming video with other animations on the same Web page. Chen is finding in Flash ways of playing with graphics, video, and sound that are unavailable in other mediums. In Sonata For The Unaware, he says he is contemplating how "music is nothing but a mess of raw frequencies forced into order by someone—a violinist or in this case a mass of Philadelphia commuters."

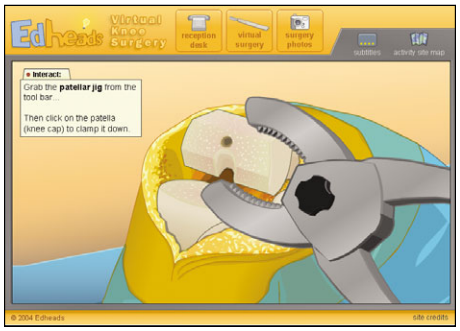

Grab the **patellar jig** from the tool bar...

Then click on the patella (knee cap) to clamp it down.

Virtual Knee Surgery • www.edheads.org/activities/knee

Call it a high-level simulation game or just online medical training, Virtual Knee Surgery hands you scalpel, drill, mallet, rongeur, and bovie (all surgeon's tools) and walks you through a knee procedure with disarming accuracy. You literally saw off a kneecap and go to the work of shaving and reconstructing a joint, all with voice-over instructions and explanations along the way. "It shows how an Internet user can be heavily engaged in a 'physical' activity like knee surgery through the ability to pick up tools and affect a virtual environment," says the project's Creative Director Eric Bort. Most of the tools available in Flash, from audio effects to interactive animation, were used for this design. "The great thing about building a project like this in Flash is that if I can think up an idea, I can find a way to make it work in Flash. The ability to do anything in one program is great."

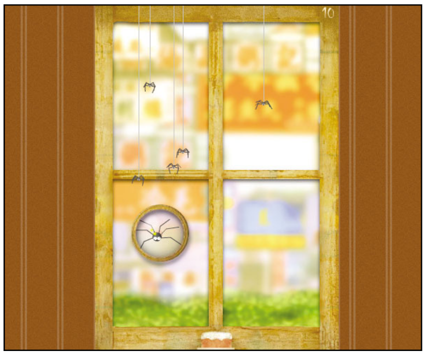

Orisinal • www.orisinal.com

There may be more elaborate Flash games and gaming hubs on the Web, but Orisinal is a truly unique collection of small programs that feel more like art than games. In "Hungry Spiders," you aim and click a magnifying glass on dropping spiders in order to protect your sandwich, but the simple arcade game looks like a folksy watercolor painting. Independent designer Ferry Halim, the freelance game developer who created Orisinal, calls it "my personal playground, a place for me to showcase my personal creations." Halim says that Flash has a number of tools that make the game developer's job much easier than other languages, and he expects much more elaborate work from programmers in the future. "Orisinal shows that Flash is really capable of making fun games. It's a great alternative tool for game development."

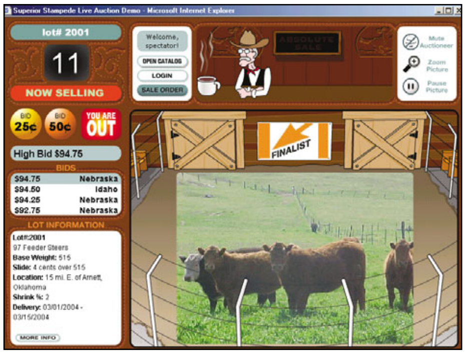

High Bid $94.75

BIDS

$94.75	Nebraska
$94.50	Idaho
$94.25	Nebraska
$92.75	Nebraska

LOT INFORMATION

Lot#:2001
97 Feeder Steers
Base Weight: 515
Slide: 4 cents over 515
Location: 15 mi. E. of Arnett, Oklahoma
Shrink %: 2
Delivery: 03/01/2004 - 03/15/2004

Superior Stampede Auction • www.stampedecattle.com

This FlashForward award winner uses Flash to run an online cattle auction, and it demonstrates exactly how sophisticated the platform has become as a front end for complex applications. Designed by Drew Falkman of Veraison and Aaron Pendland of Pendland Creative, both in Portland, Ore., the Superior Stampede Auction is a real-time auction that uses animation, a live auctioneer, and live bidding to let ranchers bid on livestock herds online. Superior Stampede sells up to 30,000 head a month to 4,000 buyers. The program makes special use of Flash's Communication Server for allowing real-time exchanges between users, but it is also combined with Java and SQL Server on the back end. For Falkman, Flash provides a unique way to bolt rich, expressive interfaces onto otherwise ordinary database and communications tools. "The interface choices are really limitless," he says.

Theban Mapping Project • www.thebanmappingproject.com

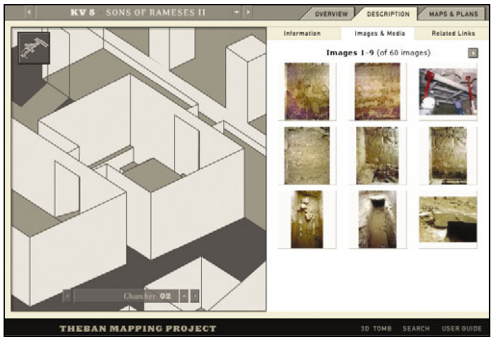

Some things are best left virtual, and many of us consider walking through labyrinthine tombs one of them. Luckily the designers at Second Story specialize in making online museums and live interactive kiosks, but the Theban Mapping Project is the mother of all educational Flash programs. The entire Valley of Kings, a burial place for 500 years of Egyptian rulers (1500 B.C. to 1000 B.C.) becomes a clickable online archeological dig. The program zooms you into the multiple pathways of the tombs and then zooms in even further to details of individual rooms, accompanied by an interactive line diagram and photos of specific areas, video, and voice-overs. All of the maps are done to scale, and the project even includes a ruler for making precise measurements within the tombs. It is an amazing trip.

Broken Saints • www.brokensaints.com

...AND FOLLOW MY LEAD.

One of the most ambitious Flash-based Web animation series is making the leap to DVD this fall as the widely hailed "Broken Saints" epic demonstrates that Flash is now ready for TV screens. It has won just about every relevant award imaginable and remains among the most creative uses of the Web for storytelling we have seen. In 24 chapters, nearly an hour long, a truly weird, apocalyptic tale unfolds involving four characters around the world whose minds are being affected by a villain's strange machine. Creator Brooke Burgess, along with Flash Programmer Ian Kirby and Art Director/Character Designer Andrew Weist, realized via Flash an animated graphic novel. The characters and background are drawn painstakingly in Japanese anime style. There are a lot of Flash cartoons on the Web, but nothing we have seen uses color, tone, character, movement, and pace in such a polished way.

Winged Sandals • www.abc.net.au/arts/wingedsandals

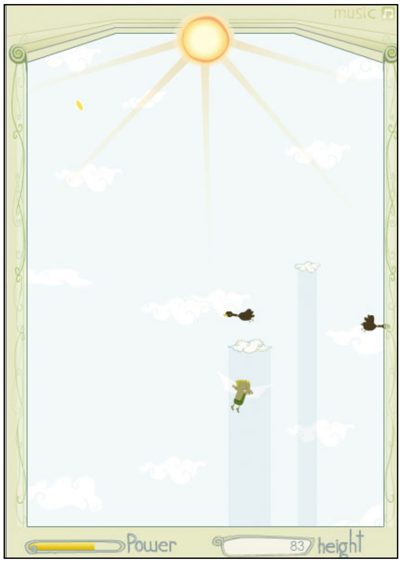

Sponsored by the Australian Broadcasting Company, Winged Sandals is a cartoonish romp through Greek mythology. The site offers animated tales of Perseus and Medusa, Orpheus and others, plus a full, clickable guide to the gods, mortals, and monsters of Greek myth, games, and craft projects all in an interface depicting ancient Greek structures that become 3D objects that zoom out to the user. According to co-producer Dee Hannan, Flash was a critical piece because the "interface is heavily integrated with a database-driven back end, which has enabled us to push information while maintaining the aesthetic integrity of the site." The heavy use of 3D models makes the whole project more immersive and kinetic. "Flash programming can be powerful when integrated with other programming languages such as HTML and ASP," says Hannan. "I think our animated character Hermes points to a new direction in freeing up some of the conventions that have developed online and actually engaging an audience with the functionality and navigation of the site so that surfing becomes enjoyable."

Red Bull Copilot: Supermoto • www.redbullcopilot.com

Some online projects are so good that it's easy to forget they're really "branded entertainment," sponsored promos for silly things like, well, Red Bull energy drinks. Red Bull Copilot from design house Science + Fiction, however, is too awesome to overlook. It pushes the video and data management capabilities of Flash so far that even Macromedia's Jeff Whatcott says "it is quite amazing." A single Flash console lets you accompany a driver in a real motorcross race, lettying you switch among seven different camera angles on the fly; monitor speed, rpms, driver heart rate, and stress level; and even switch among three different audio tracks. "It is a good example of having video content that is synchronized with graphical data displays," says Whatcott. It demonstrates how Flash can integrate multiple video streams (which switch very quickly in the program) and also connect to databases that can feed content to the viewer in real-time. What it really comes down to is how it turns a massive collection of data and video streams into cutting-edge entertainment. "You can actually see the guy's heart rate go up as he comes to an important curve," says Whatcott.

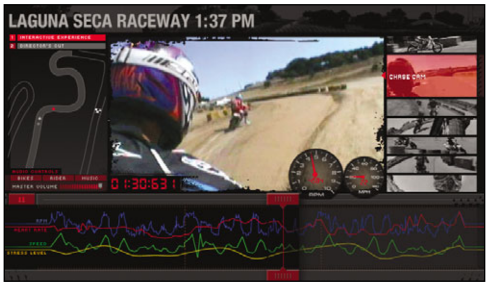

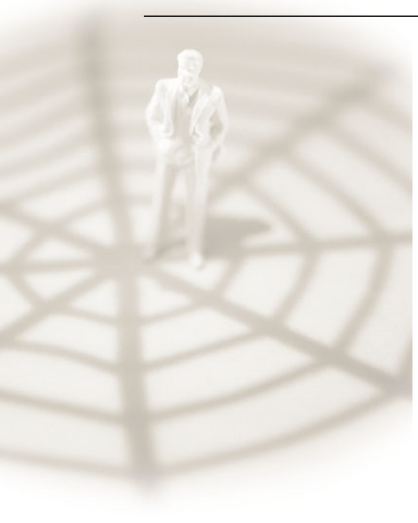

Coder's Corner: XML

XML Schema, Part II: Global & Local Types

Last month we introduced the basic model for XML Schema. Starting with a simple DTD (an external subset, stored in its own file), we introduced an equivalent XML Schema document, explaining how the schema parts work. We also discussed how an XML document can reference a schema and introduced the concepts of global and local elements. This is one of the features that makes XML Schema more flexible than DTDs. We'll look at this in more detail, as well as simple vs. complex element types.

Global Element Declarations

Global elements are declared just inside the schema's xs:schema element. The following declares three global elements: birthplace, data, and item:

```
<xs:schema xmlns:xs="http://www.w3.org/2001
   /XMLSchema" >
  <xs:element name="birthPlace" type="xs:string"/>
  <xs:element name="data">   ... complex type
   rules </xs:element>
  <xs:element name="item" > ... complex type rules
   </xs:element>
</xs:schema>
```

Rules inside a global element declaration can reference other global elements to indicate where such elements can appear. The following uses ref in <xs:element ref="item" /> to say that item elements, as defined by a global element declaration, can appear inside data elements (subject to the constraints of the complex type rules):

```
<xs:element name="data">    <xs:complexType>
  <xs:sequence maxOccurs="unbounded">
   <xs:element ref="item"/>  </xs:sequence>
</xs:complexType> xs:element>
```

Think of globals as "reusable"; once declared globally, element declarations can be referenced and used inside any complex type declaration.

A second feature of global elements is that they can appear as the root element of XML documents corresponding to that schema. Thus, the previous schema supports three potential document types, with root elements <birthplace>, <data>, or <item>. In other words, the number of global element declarations in an XML Schema is equal to the number of document types that schema supports. Consequently, if you don't want an element to appear as the root element, it shouldn't be defined as a global.

Global Attribute Declarations

Global element declarations are equivalent to <!ELEMENT declarations in DTDs. But XML Schema also supports global attribute declarations, which isn't possible with DTDs. For example, the following declares a global attribute named id:

```
<xs:schema xmlns:xs="http://www.w3.org/2001/
   XMLSchema" >
  <xs:attribute name="id" type="xs:ID"
   use="required" / > ...  more global types ...
</xs:schema>
```

Using this global, we can write the declaration (introduced last month) for item elements as:

```
<xs:element name="item">
<xs:complexType> <xs:sequence> <xs:element
   ref="name"/> ... other elements
   ....</xs:sequence>
<xs:attribute ref="id" /> </xs:complexType>
</xs:element>
```

Global attributes aren't as common as global elements, but they can make it easier to manage a schema. One example is ensuring consistency of attribute types across different element declarations. Consider a large schema with id attributes associated with hundreds of elements. Declaring id to be global and then referencing this global definition means such xs:attributes will take the same typed values.

XML Schema also supports local element and attribute declarations, or declarations for elements or attributes that appear inside an element declaration and that apply only in that local context. The following declaration for item elements is equivalent to the previous one but uses local declarations for the name element and id attribute:

```
<xs:element name="item">
<xs:complexType> <xs:sequence> <xs:element
   name="name" type="xs:string" /> ... other
   elements ....</xs:sequence>
```

Coder's Corner

```
<xs:attribute name="id" type="xs:ID"
   use="required" /> </xs:complexType>
<xs:element>
```

Using local and global declarations isn't an either/or decision. A schema can use both types at the same time. For example, a schema can define a global comments element to have a particular content model but elsewhere define a local comments element to have completely different allowed content. Both global and local declarations for elements with the same name can be referenced and used inside the same definition. The following schema has a global declaration for comments elements and uses this and a local declaration for comments inside the item element declaration:

```
<xs:schema
   xmlns:xs="http://www.w3.org/2001/
   XMLSchema" >
 <xs:element name="comments">
   <xs:complexType mixed="true">
     <xs:choice minOccurs="0"
   maxOccurs="unbounded">
       <xs:element name="para"
   type="xs:string"/>
       <xs:element name="heading"
   type="xs:string"/>
     </xs:choice>
   </xs:complexType>
 </xs:element>
 <xs:element name="item">
   <xs:complexType>
     <xs:sequence>
       <xs:element name="name"
   type="xs:string" />
       <xs:element name="comments"
   type="xs:string" minOccurs="0"/>
       <xs:element ref="comments"
   minOccurs="1" />
     </xs:sequence>
   </xs:complexType>
 </xs:element>
</xs:schema>
```

The global declaration for comments says this element can contain numerous para or heading elements (which in turn contain text), interspersed with text (this is the meaning of the mixed="true" attribute on the complexType element). The global declaration for item says this element contains a sequence of three elements: a name element, followed by the locally defined rule for comments elements, followed by a comments element consistent with the global declaration. Here's an example:

```
<item>
 <name> Ian Graham </name>
 <comments> some text comments --
   note this element can be absent.
   </comments>
 <comments> <heading> About Ian
   </heading> <para> Ian writes the
   Coder's Corner articles for ... </para>
   </comments>
</item>
```

Avoid the temptation to declare an element with two attributes having the same name, one local and the other global. XML syntax disallows this by saying that an element can have only one attribute with a given name. Thus, the following markup is syntactically incorrect, independent of any schema or DTD:

```
<errorElem attr="value1" attr="value2" />
```

Simple vs. Complex Types

XML supports two models for element or attribute types: simple and complex. Simple types are constrained text strings with no markup or other structures, a string ("hello there!") or a real number (1.23e+23) being examples. Attributes can only take simple type values; XML syntax rules say that attributes can only have string values. XML Schema predefines 44 simple types, such as string, ID, etc. Declarations such as the following simply declare the element to be one of these predefined types:

```
<xs:element name="name"
   type="xs:string" />
```

Complex types correspond to values that have structure, such as text strings combined with markup elements or elements having attributes. Thus elements can be simple or complex. In short, an element is a simple type if it contains just a string of text (but no markup), and if it has no attributes. Conversely, an element is a complex type if it has attributes, if it contains markup in addition to text or both. There are no predefined complex types. Most of the job of designing a schema is in creating the needed complex types. Complex type rules are defined inside an xs:complexType element.

XML Schema supports many attributes and elements to define complex type rules. The example below shows how, inside an xs:complexType, an xs:sequence element defines an ordered sequence of allowed elements and how attributes (xs:attribute) are attached to the element. It also shows how minOccurs/maxOccurs (with defaults of 1) attributes on xs:element can specify how many of these elements can appear at this place in the sequence:

```
<xs:element name="item">
   <xs:complexType>
<xs:sequence> .. sequence of element declarations ... <xs:element ref="comments"
   minOccurs="0"/>
</xs:sequence> <xs:attribute name="id"
   type="xs:ID" use="required"/>
         </xs:complexType> </xs:element>
```

This example introduces a new property (mixed="true") for complex types, and an xs:choice element:

```
<xs:complexType mixed="true">
<xs:choice minOccurs="0"
   maxOccurs="unbounded">
<xs:element name="para"
   type="xs:string"/>
<xs:element name="heading"
   type="xs:string"/>
</xs:choice></xs:complexType></xs:
   element>
```

Mixed="true" means the complex type can intersperse text among the components defined inside the xs:complexType (xs:choice and xs:sequence being component examples). The xs:choice element defines a block, but here there's a choice of one of the elements defined inside (para or heading). MinOccurs/maxOccurs define how many times the chosen element can appear inside the block. The allowed content inside comments elements can be a mixt of text, para, and heading elements in any order and with any number of successive paras or headings.

Next month we'll look more at complex types and mechanisms for defining types independent of elements.

by Ian Graham

(For more on the documents in this article, go to www.utoronto.ca/ian/articles/aug04.)

Caught In The Web

by Rob "Cmdr Taco" Malda

bizarro.txt

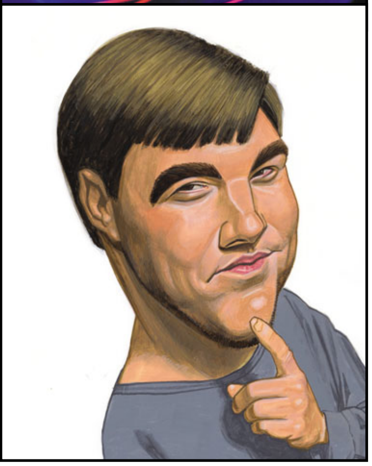

Rob "CmdrTaco" Malda is the creator and director of the popular News for Nerds Web site Slashdot.org. He spends his time fiddling with electronic gizmos, wandering the 'Net, watching anime, and trying to think of clever lies to put in his bio so that he seems cooler than he actually is.

They aren't new, but like so many things, the Internet has made what I'm calling "Bizarro Fans" an even more vocal part of the media. Today they affect any conceivable area of interest that has a voice online, be it politics, technology, or any piece of popular or unpopular culture. These users add immensely to the discussion of any subject, and at the same time, they drag it down.

The Bizarro Fan is amplified by the Internet. The vast majority of people, upon encountering a subject that they don't like, would simply ignore it and move on. Usually, the ranks of the Bizarro Fans come from the larger body of actual fans. People who like the Web site, TV show, politician, or whatever the topic is. This means they are already personally invested in some way in the subject. They feel that they are owed something, so they react.

I choose to ignore them wherever possible.

Bizarro Fans lash out in many ways, and the Internet truly has given them voice. Decades ago, a fan could send a postcard to his favorite actor, care of the TV show she starred in. It was less likely for the Bizarro Fan to go to this length, and because the fan's actions rarely had any visibility, the problem remained relatively small. But on the Internet, the Bizarro Fan usually sees the true fans engaging in their happy activities relating to the subject at hand. This happiness enrages the Bizarro Fan, and like a child, he lashes out at the community.

The ongoing battle between fan and Bizarro Fan has been played out online in countless ways in the short life of the Internet, and it's important for everyone involved to know how to recognize these people. Most experienced users have learned to ignore them, but others will fall right into the trap. They will engage the Bizarro Fan in a useless repetitive debate. Both sides have probably been through the discussion before, but it is still played out time and time again, fueled entirely by the rage of the Bizarro Fan with time to spare and a forum to pollute. The fan's anger is fed by an audience and amplified by anyone willing to enter into the irrelevant debate.

This can be really disheartening for the fans and for the creative folks who are responsible for the subject of all this controversy in the first place. But let me stand up and tell you that the Bizarro Fan is a completely normal occurrence online. In fact, not only is it normal, it should be considered a form of flattery. It is the most powerful indicator that you and your site have succeeded in making a difference: The Bizarro Fan is your best indicator of actual fans.

Some of the most common questions posed to me revolve around these types of people and how Webmasters ought best to deal with them. I think the first thing that any aspiring Webmaster needs to understand is that these users are the single greatest indicator of success. Simply gaining traffic is a reasonable measure of interest, but not a measure of passion. When you have risen to the point where there are Internet users out there with this level of hostility about your work, you can be assured that the work you are doing matters, no matter what the subject.

As for dealing with Bizarro Fans, you have many choices. Personally I choose to ignore them wherever possible. As long as their actions don't interfere with the participation of other users, it is best to let them rant and rave. Any involvement on your part will simply feed their rage. These are people who want to be angry—whatever you do ultimately feeds that emotion. They will never go away, but you don't need to give them an incentive.

When dealing with them on your own forums, you can choose to invent some sort of moderation system (or use www.slashcode.com, the open-source application that powers Slashdot). Our system tries to hide their debates from average users but still make them available to those who want them. Other sites simply delete posts that they feel interfere with their forum, but I don't think that this is a very scalable solution, to say nothing of the fact that having their words deleted simply fuels Bizzaro Fans' anger and gives them one more thing to complain about.

Ultimately you need to accept that these people are an annoying but normal part of any community. Accept them as proof that your work matters. Accept the fact that they are an incredibly vocal minority. Accept the fact that their actions unintentionally are proving that you are doing something right. ■

Rob ignores any bizarro criticism sent to malda@cpumag.com.

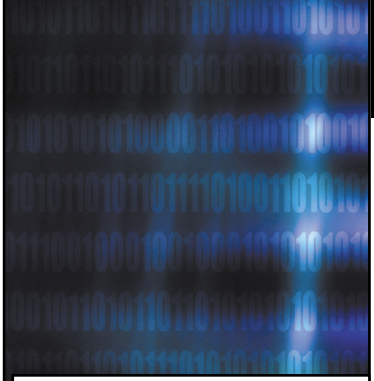

Forward Slash

by Joan Wood

Row, Row, Row Your Boat

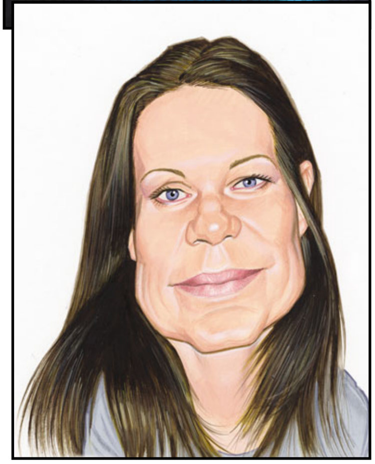

After a decade of video engineering, Joan migrated to game development, co-founding Xatrix Entertainment and producing the two Cyberia titles. As president of indy developer Mango Grits, she talked 3Dfx out of prerelease Voodoo hardware and developed flying game Barrage for Activision, then co-founded and ran hardware review site SharkyExtreme.com as managing editor. Ensconced in Silicon Valley, Joan currently plays with small, wearable PCs and big, rack-mount vis-sim image generators while secretly plotting to save the world through a series of international location-aware multiplayer handheld games.

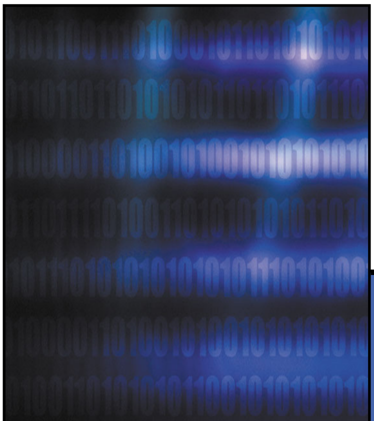

As SCO has made headlines attempting to make a mint while thwarting both Linux and the concept of open source, another patent battle has been quietly taking place that puts a kink in all things streaming, from satellite cable broadcasting, to pay per view, to Internet radio. Even video on demand in the privacy of your own seedy motel room is no longer safe. Apparently, streaming media has become as ubiquitous (and patentable) as 1-Click shopping.

If you've never heard of Acacia Research Corporation, be prepared for a revelation. The company claims that its acquired DMT (digital media transmission) patents cover transmission and reception of stored digital content over Internet, cable, satellite, wireless, and other unspecified delivery methods. Acacia has spent the last year or so persuading (sometimes with the help of lawsuits) smaller users of streaming technology, primarily adult entertainment Web sites and hotel in-room video-on-demand services, to sign licensing agreements. It has now moved up to the big time, taking on major cable and satellite TV companies and making noises at wireless telcos delivering mobile media.

Streaming media has become as ubiquitous (and patentable) as 1-Click shopping.

With the filing of a patent infringement lawsuit against some smaller players, along with giants Comcast, Cox, Charter, Echostar, and DirecTV, Acacia appears to be betting that the groundwork laid through the earlier license agreements (over a hundred of them) in favor of its five DMT patents will be persuasive enough either to leverage more license agreements or obtain actual judgments against the telecommunications companies. Signing a license deal with the Walt Disney Company probably didn't hurt.

Who'll Stop The Stream?

Don't get me wrong, enforceable patents (and caffeine) form the backbone of modern invention, and I personally wouldn't have it otherwise. Commercial streaming media users have known about this potential bombshell for a while now, but there was always a hint that it might just go away; that Acacia's claims were too broad-reaching to be enforceable. There was an idea that when things got really out of hand, when Acacia went after the big patent violators like Microsoft, Apple, and RealNetworks, there would be a push back. Not so. Acacia wisely labels MediaPlayer and RealPlayer as beneficial contributory infringers and instead goes after the media companies making money directly from streaming through those players. By targeting violators with royalties tied to a percentage of overall business revenue, companies have an incentive to settle sooner rather than later to get good royalty rates per download or per customer instead.

What else is out there waiting? Acacia Research Corporation seems to follow the trend of being more of an IP licensing company than a true R & D entity. As its coffers fill with streaming media and V-Chip (another patent it has acquired) money, the company may buy up other general-purpose process patents from near-death dot-coms. By creating a portfolio of IP and patents in areas where it does no actual work, Acacia can avoid both contributing to evolving technology standards and the risk of the patent infringement counterclaims that keep chip company lawyers happily engaged in negotiation-through-lawsuit cross-licensing agreements. Not invention's finest hour.

The Road Not Taken

Just to end on a more positive note, a reminder of a bifurcation milestone we should all hold dear. When things look bleak, perhaps it would help to ask ourselves where we might be had Tim Berners-Lee and other Web pioneers decided to make their fortunes off the patentable invention they created instead of insisting on license-free technology. Even more remarkable was the decision by his employer, CERN, not to pursue royalties, either. Just try running those numbers through a dead dot-com and see how many venture capitalists sign up. The last laugh has to go to Berners-Lee, who has been awarded the Finnish Technology Award Foundation's Millennium Technology Prize. Supported by the Finnish government and private contributors, it is valued at 1 million euros ($1.2 million). Sometimes good things happen to good inventors. ■

Invent comments and send them to joan@cpumag.com.

by Chad Denton

Road Warrior

Broadband Over Television, Free Ringtones, Motion-Sensitive Nokia Accessory & More From The Mobile Front

Free Ringtones For All?

Wireless customers have enthusiastically embraced ringtones. Of course, when there's money to make, what's popular with consumers quickly becomes popular with providers. Depending on the service provider and ringtone, a 30-second ringtone can cost anywhere from $1 to more than $7. In some cases users can't even keep the ringtone.

Sprint PCS users who spend $1.50 to $2.50 per ringtone, for example, can only keep it 90 days. For 99 cents you can buy a song on iTunes and keep it. Some stereo ringtones that sound close to CD quality are going for as much as $7, due mainly to licensing fees music companies demand. Music publishers and musicians aren't entitled to as much of a cut from polyphonic, monophonic, and MIDI tones.

New software from Xingtone, however, might put a crimp in the ringtone market. For less than $15, Xingtone lets users create an unlimited number of ringtones from their music collection and download the tunes to their phone.

Predictably, some music executives are upset with a product that may usurp a developing source of revenue. Xingtone, however, doesn't have all within the music industry petrified. The band Sugarcult, for instance, includes a limited version of Xingtone with its latest CD "Palm Trees And Power Lines." The stripped down version lets fans make ringtones from tracks on the CD. Wireless providers are also concerned about a potential loss in revenue, although some feel Xingtone may be a way to promote their wireless services.

Xingtone software ($14.95) isn't compatible with all phones or services. A trial version lets users test the software with their phone using prepackaged ringtones. (You'll have to pay up before you can create your own ringtones.) Some services, such as Nextel and Verizon Wireless, aren't compatible with Xingtone because they don't let users download unapproved content. ▲

Motorola's V600 from AT&T Wireless is one of 52 cell phones Xingtone is compatible with. You can try out your phone by downloading the free demo version of the software.

Cell Phones For Soldiers

Just because you're young doesn't mean you can't make a difference. Two Maryland middle school-aged sisters, Brittany and Robbie Bergquist, collected $16 in early April in an effort to help Sgt. Bryan Fletcher pay the $3,411 in cell phone charges he's amassed while serving in Iraq. Although T-Mobile eventually forgave the bill, the girls continued raising money and recently topped the $100,000 mark. As a result, 400 members of the Massachusetts National Guard Battery 42nd Division Artillery received a 200-minute prepaid calling card when they arrived in the Middle East last May. You can contribute to the cause by sending a donation to Cellphones for Soldiers; c/o South Shore Savings Bank; 5 Assinippi Ave; Hanover, MA 02339. Alternatively, you can donate any new or used cell phones to Cellphones For Soldiers; c/o Standard Communications Group; 31 Schoosett St., Suite 205; Pembroke, MA 02359. Visit www.cellphonesforsoldiers.com for more information. ▲

Nokia 3220 Xpress-on Adds Light Messages

In terms of style, Nokia might be the Apple of the mobile phone world. Nokia plans to introduce its new 3220 phone in Europe, Asia, and yes, even North America. The phone includes a camera, photo-editing software, games, and customizable ringtones and wallpapers. The phone's back is transparent, letting users insert cutouts to personalize their phones.

Of course, none of this is really new. What's interesting is the optional Xpress-on Fun Shell that Nokia is selling for the 3220. The Fun Shell includes LEDs on the shell's back that lets the user write in midair. Users can choose from a set of preprogrammed messages or create their own messages. As they wave the phone from side to side, the LEDs flash to make a message appear in midair. The Fun Shell also includes two games that take advantage of the Fun Shell's motion sensors. Prices weren't available at press time. ▲

Nokia's 3220 should be coming to North America this summer and will feature an optional shell that uses LEDs to write messages in midair.

Wireless Broadband Over Television

In early May the FCC issued a Notice Of Proposed Rule Making that would allow for unlicensed data broadcasts over the unused television spectrum between channels 5 and 51. No, "unlicensed" isn't a mistake. Traditionally, television airwaves have been some of the most sought after because lower frequencies can travel farther and penetrate walls and other obstacles more easily. Allowing unlicensed use of this spectrum means anyone that adheres to FCC restrictions could use this low-frequency spectrum.

FCC restrictions would include the use of so-called "smart radio" technology to avoid interference with existing television stations. Such devices would utilize unused channels for broadcasting data. Naturally, the National Association Of Broadcasters isn't too pleased with this, believing the potential for interference still exists.

The FCC envisions the rule giving birth to low-powered, short-range, wireless devices (similar to existing Wi-Fi products) and high-powered fixed wireless stations that could deliver broadband to entire communities. At this point the rule change is only a proposal, but it could be a move in the right direction for wireless users. Besides, does anyone still get TV over the air? ▲

Wireless Number Database Moves Forward

In February 2004's "Road Warrior" (page 83), we reported that the CTIA is developing a national database for wireless phone numbers. At the time, the CTIA was pushing for an opt-out database, but it now insists the database will be opt-in. The CTIA continues to say it won't publish the database in paper form or sell the data to third parties.

The database will be a 411-style directory, with a small fee charged for using it. The CTIA says the database could be beneficial to real estate agents and other mobile professionals who want their numbers readily available but currently have to pay extra to list their numbers.

Will the database make it possible for telemarketers to interrupt your dinner, even when you're at a restaurant? Current FCC rules prohibit autodialed and artificial or prerecorded messages to wireless numbers. The rules, however, don't prohibit live agent calls to your wireless number. If the CTIA sticks to its promises, however, telemarketers would have to pay a 411 charge for each number it collects from the database. This could make compiling such a list potentially expensive and inefficient. If you're still worried about telemarketing calls, you can register your number with the National Do Not Call List (www.donotcall.gov).

If spammers gain access to the database, SMS spam could be a bigger problem. The FCC doesn't allow SMS spam, but most spammers regard such laws with the same type of respect that some of us regard processor warranties with. If the CTIA lives up to its promises, it would be difficult for spammers to get your number, but anything is possible.

U.S. Representatives Joe Pitts (R-Pa.) and Edward Markey (D-Mass.) are working to get their Wireless 411 Privacy Act through Congress before the database goes live. The bill would require existing users to opt in and prevent providers from charging extra fees for unlisted numbers.

Not all wireless carriers have signed on to the database. A major exception is the largest wireless carrier in the country, Verizon Wireless, which will continue to charge a small fee for customers who want their numbers published, but it doesn't plan to provide any additional numbers to the CTIA database. ▲

Dell's New Pocket PCs

It seems there may be a bit of a price war going on in the PDA market. Last month we mentioned how impressed we were with the prices of new Palm models, especially the Zire 72. Now Dell's new Axim X30 Pocket PC line also has some impressive features, in addition to an impressive price tag.

The entry-level X30 includes a new 312MHz Intel XScale processor, along with 32MB of RAM and 32MB of Intel's StrataFlash ROM. The entry-level model sells for just $199. For an extra $50, you can get a model with double the memory, plus integrated Bluetooth. The $349 model (just $50 more than the Zire 72) includes a 624MHz Intel XScale processor, 64MB of RAM and ROM, integrated Bluetooth, and integrated 802.11b. All models include an SDIO-compatible expansion slot, a removable rechargeable battery, and a 3.5-inch, 16-bit color display.

All three models run Microsoft's Windows Mobile 2003 Second Edition. The new OS is more of a refresh than a major update. Changes include a one-column view in Internet Explorer and support for switching the display to landscape mode. Pixel-doubling will let you view any application in landscape mode, even if the application doesn't include native landscape support. ▲

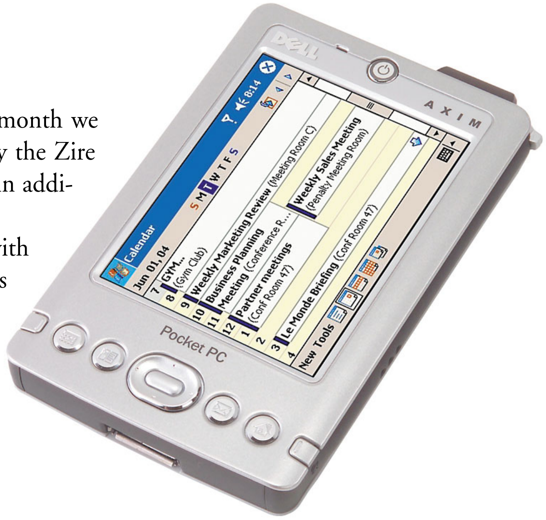

The high-end Axim X30 offers Bluetooth, Wi-Fi, and a 624MHz XScale processor for just $349, while the low-end model offers a 312MHz XScale processor for $199.

Compiled by Samit G. Choudhuri & Chris Trumble

At Your Leisure

Plug In, Sit Back & Fire Away

The entertainment world, at least where it pertains to technology, morphs, twists, turns, and fires so fast it's hard to keep up. But that's exactly why we love it. For the lowdown on the latest in PC entertainment, DVDs, consoles, and just stuff we love, read on.

Thief: Deadly Shadows

The Original Sneaker Is Back

Those familiar with the previous Thief games will recognize the return of master thief Garrett. You'll once more be skulking around medieval manors stealing from the rich to give to yourself. Your trusty blackjack will serve you well as you explore your surroundings and uncover the intriguing storyline that often progresses as you eavesdrop on conversations. What you won't recognize is the third-person camera—relax, you can switch between third- and first-person views with ease. Third-person gives you a better feeling for your environment, but we stuck to the first-person mode for its more "I'm there" feeling.

The dynamic lighting is well-suited to the gameplay and is an effectively immersive tool. But characters move around like animatrons and are just not *convincing*. Frame rates were spotty on both Xbox and mid-range PCs but good on the highest-end rigs. The AI was mostly OK but had bouts of incredible stupidity at the Normal skill level. Kicking up the difficulty made the game more interesting. That said, we should also note that the Xbox version has a bug that loads saved games up at Normal difficulty regardless of your preset difficulty level, which nullifies our earlier difficulty level suggestion. No fix was being offered as we went to print.

Thief: Deadly Shadows is not a breakthrough title but will provide ample enjoyment for fans of the genre. And as things stand with the Xbox bug, the PC version is definitely the way to go. ▲

ESRB: (M)ature • Eidos Interactive

$39.99 (PC) • $49.99 (Xbox)

www.thief3.com

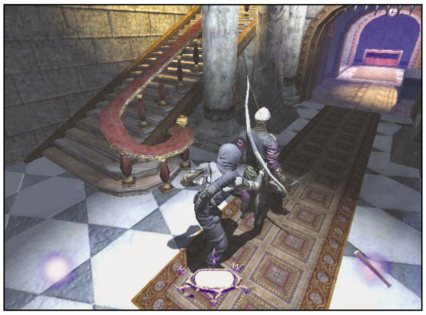

You can sneak up behind enemies with a dagger to finish them off. Feeling kind? Use your blackjack to knock 'em out cold instead.

DVD Byte by Todd Doogan

THE DIGITAL BITS
we know dvd
www.thedigitalbits.com

July brings us a sense of patriotism and a nice selection of DVDs to burn the time away. This month doesn't offer the selection we'll have in the fall, but it's not too bad. July 6 brings the first volume of **Batman: The Animated Series**, Ashton Kutcher in **The Butterfly Effect**, the Aussie rock film **Garage Days**, and classic superheroes prancing around in the first collection of H-B's **Challenge Of The SuperFriends**. July 13 sends us Meg Ryan as a boxing promoter in **Against The Ropes**, a special edition of **The Bourne Identity**, the steamy NC-17 cut of **The Dreamers**, and **Zatoichi Meets The One-Armed Man**. Fans of Adult Swim will be happy on July 20 with the second volume of **Aqua Teen Hunger Force** and the first volume of **Sealab 2021**. Lindsay Lohan fans will want **Confessions Of A Teenage Drama Queen** in their collection, and Arnie fans will be happy with a new Collector's Edition of **Predator**. And don't forget **Millennium** and both the TV version and the new film version of **Starsky And Hutch**. The month comes to an end with a one-two punch from Guillermo del Toro with special editions of **The Devil's Backbone** and **Hellboy** on July 27. Also look for the classic anime **Neon Genesis Evangelion Platinum Volume 1**, **Showgirls: The VIP Edition**, **Sledgehammer: Season One**, and **The Whole Ten Yards**. Now you have something to do while you eat your hotdogs and potato salad. ▲

Driv3r

A Hyped-Up Disappointment

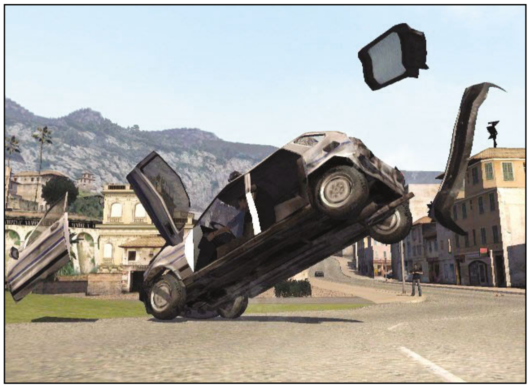

Admittedly, there are shots where the graphics look good, but a deeper look reveals jaggies. The biggest problem, though, is that Driv3r isn't really all that much fun to play.

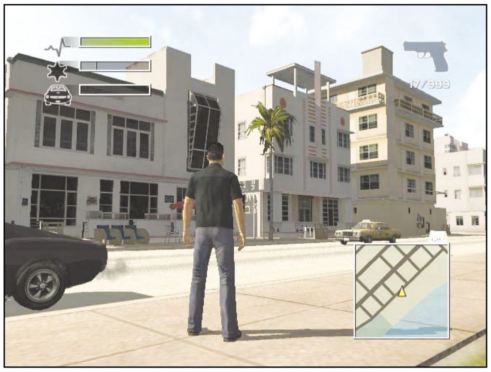

Tanner ponders if a game of Tiddlywinks or Badminton might be a better way to spend his time.

In Driv3r, you step into the shoes of Tanner (voice by Michael Madsen) and go undercover to infiltrate a Miami gang that's stealing cars. In your quest to get answers, you'll travel from Miami to Nice (France) to Istanbul (Turkey).

Voice actors whose names you might also recognize are Ving Rhames, Michelle Rodriguez, Iggy Pop, and Mickey Rourke. Unlike Grand Theft Auto, these thespians are introduced in the opening credits so you won't miss noting their participation.

Driv3r had received a lot of hype from Atari, so we were disappointed to discover a very mediocre game under a light veneer of polish. Production values are top-notch with a good soundtrack, but it's a bad sign when cut scenes (excellent as they may be) are the main highlight. The on-foot control feels clumsy and limited, but vehicular control is better (though not great). Plus, the wacky physics thoroughly ruin your suspension of disbelief—just wait until you see a car spinning like a top while staying stationary in a single location *midair*. It doesn't happen all the time, but you get our issues with the physics–it ruins the suspension of disbelief. We were only able to play the Xbox release but will let you know if the PC version is an improvement when it comes out in October. ▲

ESRB: (M)ature • Atari

$49.99 • Xbox, PS2

www.driv3r.com

Ultrasone HFI-15G Headphones

Surround Gaming Audio Done With Style

These headphones are comfortable and lightweight. Most importantly, they sound excellent.

We play our Xbox and PS2 games in Dolby Digital 5.1 and PC games with a Sound Blaster Audigy 2 ZS hooked to 5.1 speakers. However, it's not always possible to play games or watch DVDs in full surround sound at 3 a.m. Apparently there are sleeping individuals who can be awoken by directional audio effects and loaded soundtracks.

To that end, we recently set out to find surround-sound headphones to use during our late-night gaming sessions. The current crop of 5.1 headphones just didn't make the grade in terms of quality and cost, and we were left searching for higher quality at lower prices. We finally happened upon a set of headphones (the HFI-15G) by a German company called Ultrasone. These headphones are available in two models: LE and PROline. They sound identical, but the former shields up to 60% of magnetic field emissions, while the latter apparently takes care of up to 95%. The HFI-15Gs have a frequency range of 20Hz to 20KHz with an im-pedance of 32 ohms, so you'll be okay powering these headphones with a mobile player.

We tested the 3.85-ounce PROline headphones (the LEs are .875 ounces lighter) extensively with Thief: Deadly Shadows, Splinter Cell Pandora Tomorrow, The Chronicles Of Riddick: Escape From Butcher Bay, and several other games. The HFI-15Gs performed admirably in all our tests over many, many hours. The surround-sound audio using the headphones' built-in S-Logic feature never led us astray—positional cues were dead-on and let us reap the benefits we also get from our 5.1 speakers. Best of all, the headphones were comfortable, sounded clean, and minimized the radiation to our noggins. Give these a go if you need quality surround-sound headphones for LAN parties and late-night gaming sessions. We came away impressed. ▲

Ultrasone

$99.99 HFI-15G LE; $149.99 PROline

www.ultrasoneusa.com

The Chronicles Of Riddick: Escape From Butcher Bay
What? Another Good Movie Game?!

Say hello to my little friend.

Last month we reviewed the PS2 and Xbox game Van Helsing, a movie-licensed joint that was far more fun than its silver-screen namesake. It looks like a pattern may be forming here: we took a look at The Chronicles Of Riddick: Escape From Butcher Bay for this issue, and it's better than the recent Riddick movie tie-in, as well.

TCOR is a first-person shooter with some interesting action-adventure elements that recounts the titular hero's escape from a maximum-security penal facility called Butcher Bay. Butcher Bay is located on a desolate desert planet, so escape isn't much of an issue—until Riddick shows up. From the moment the game begins, you'll guide his efforts to get out of the hellish facility.

Get a load of the incredible detail and fantastic lighting effects you'll enjoy in TCOR: EFBB.

TCOR is special in two important ways. First, it is the best-looking Xbox game released to date. One or two titles this fall will rival its impressive visuals, but for now, second place isn't even close. Incredible lighting effects and a new rendering technique called normal mapping, which creates 3D models that look incredibly complex but are easy on the Xbox's silicon, combine to give the game a visual style that you'll currently only find in leading-edge PC games played on high-end systems. (Speaking of which, we'd love to see a PC version of this game.)

Second, TCOR is one of the most successful attempts yet to blend several game genres and elements into one package. TCOR is at its core a shooter, but it also incorporates compelling hand-to-hand and melee weapon combat and is the most natural combination of stealth and action going. And although the camera changes from first-person to third-person in certain, frequent circumstances, the switch feels natural and not at all disorienting. On top of all that, the game leans heavily on in-game dialogue that is enjoyable, thanks partly to outstanding voice talent and in part because players frequently have several dialogue options when responding to NPCs, which creates a greater sense of involvement with the story.

The game's sound effects and music are spot-on, as well. Aside from a few rather minor control complaints, there's nothing about The Chronicles Of Riddick we don't like. It's a gritty, highly atmospheric and engaging game that fires on all cylinders. We recommend it without reservation. ▲

ESRB: (M)ature • Vivendi Universal Games
$49.99 • Xbox
www.riddickgame.com/us

Front Mission 4
Final Fantasy Tactics With Mechs

Front Mission 4 plays quite a bit like another great Square game, Final Fantasy Tactics, except the units you control are giant mechs called Wanzers instead of fighters and wizards. Every turn, you move your mechs and attack enemy units or repair friendlies while pursuing basic mission objectives (destroying all enemies or defending a position or base).

There's a lot of dialogue to wade through, but there are also plenty of cut scenes that keep the somewhat complex story moving while living up to Square's rep for great CG. And, of course, you'll spend time tweaking your Wanzer pilots' skills and customizing your hardware.

One of our favorite things about FM4 is its Link System, a new strategic element that rewards planning and sound tactics. Using the Link System, your Wanzer pilots can string extra attacks together when in proper formation and do much more damage to enemies than they otherwise would.

The game's long, turn-based battles and involved stat system probably won't appeal to action junkies. But for people who enjoy winning with tactics as much as winning thanks to quick reflexes, Front Mission 4 is as fun and addictive as any console strategy game we've played. ▲

ESRB: (T)een • Square Enix
$49.99 • PS2
www.square-enix-usa.com

Susan will make sure you're comfy with your Anthro choice.

Yes, life can be this good. You, your laptop, and a comfy place to work (that's the part Anthro supplies). What more could you need? All Anthro furniture is modular and flexible, so we're delighted to help find the right fit for you. Tell us about your furniture challenges and we'll configure choices for you to consider. (Did we mention our Lifetime Warranty?) Contact our helpful folks at **Anthro.com or 800-325-3841.**

ANTHRO®
TECHNOLOGY FURNITURE®
-Fast and Friendly!

GSA contract. Available for OEM applications.

Destination Dallas

CPU GOES TO THE CROSSROADS GUITAR FESTIVAL

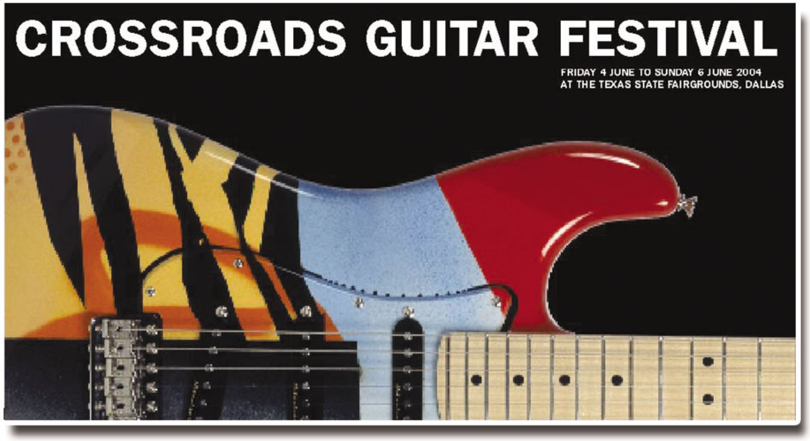
FRIDAY 4 JUNE TO SUNDAY 6 JUNE 2004
AT THE TEXAS STATE FAIRGROUNDS, DALLAS

> **I went to the crossroads,**
> **fell down on my knees**
> **Asked the Lord above,**
> **have mercy now**
>
> **- Robert Johnson,**
> **“Crossroads Blues”**

Legend has it that Robert Johnson sold his soul to the devil at straight-up midnight sometime in 1930 at the crossroads where Highways 61 and 49 intersect outside Clarksdale, Miss. In exchange, the devil gave Johnson otherworldly six-string skills. Johnson died just years later in 1938. He recorded just 29 songs in his short 27 years, but those tracks and Johnson’s insane chops remain as strong an influence on guitar players today as ever. Take Eric Clapton, for example, a lifelong Johnson devotee. Earlier this year, Clapton recorded several Johnson tunes for his album “Me and Mr. Johnson.” Clapton subsequently organized and hosted the Crossroads Guitar Festival on June 5 and 6 in Dallas, and *CPU* was there. Why? Well, for the music, of course, but also because they were using cool technology to record the show, and we wanted a look.

Eric Clapton in an impromptu Saturday jam on the Esplanade.

Roadtrippin’ With A Mission

We aren’t morning people around here, but when the chance comes to see B.B. King, Buddy Guy, Carlos Santana, Jeff Beck, ZZ Top, James Taylor, Steve Vai, Bo Diddley, John Mayer, Joe Walsh, and others, you get your bones outa’ bed at whatever time it takes, even if that’s 3:30 a.m. to catch a plane. Likewise, when you record a show of this magnitude—watch for an upcoming DVD and PBS special soon—you don’t hire some hack for the job, and you don’t use second-rate equipment. That’s why Warner Brothers Music turned to multiple Grammy-winning engineers/producers/mixers Elliot Scheiner and Ed Cherney. In turn, Scheiner and Cherney used AMD dual-Opteron-based systems to capture the music.

Scheiner and Cherney weren’t household names to us, but we were aware of the results of their work, which include efforts by Steely Dan, Bob Dylan, The Rolling Stones, Beck, The Flaming Lips, The Eagles, REM, Aerosmith, Fleetwood Mac, Billy Joel, Bette Midler, Elton John, Van Morrison, Jimmy Buffett, and others. Throw in a couple dozen or so Hollywood soundtracks, and we’re talking heavyweights here. Some further investigation led us to learn about Scheiner’s pioneering work in 5.1 Surround Sound mixing, where he’s considered the godfather. Ever hear “The Eagles—Hell Freezes Over” in DTS? Quality-wise, it’s universally considered the best live surround mix ever. That’s Scheiner’s doing. Queen’s “Night at The Opera”? That’s Scheiner, too. And check out the eight-speaker Panasonic 5.1 Surround Sound ELS system in the Honda Acura TL (www.acura.com). It’s the first North American car with such a system, and you can thank Scheiner for helping design it.

Scheiner spends much of his time today remixing classic albums in 5.1 for DVD-Audio rereleases, for which he’s somewhat of an industry guru. Needless to say, when the opportunity to meet Scheiner in Dallas to talk audio, cars, and technology came up, we were more than intrigued. When we discovered that Scheiner and Cherney would be using an AMD dual-Opteron-based DAW64 workstation to record the marathon concert (the first time the technology was used for a live concert), we were convinced Texas was where we needed to be. (Check out the “AMD Dual Opterons For The Digital Audio Generation” sidebar for more on the DAW64s.)

This dual Opteron box sat in Elliott Scheiner's trailer.

Now you also know what an analog mixer looks like.

Scheiner gives CPU a tour of his mobile recording studio.

Weapon Of Choice

We met Scheiner the day before the show outside the Cotton Bowl on the Texas Fairgrounds. Despite being a studio rat since 1967, Scheiner hasn't exactly always been down with technology. He says he "tolerated computers" until around 1996 when he knew he "needed to go digital, but I backed away from it for years." Scheiner says he realized digital was "where it was going, and I had to do it or I wasn't going to work anymore."

Early on Scheiner worked on a Mac, but he says he spent "more time rebooting" than working and was "ready to sell it in five days." After transitioning through various systems, he moved to the DAW64 running Pinnacle Steinberg Nuendo software after AMD's Charlie Boswell introduced him to the system. (See our interview with Boswell on page 103.) "It was mind blowing," Scheiner says.

"It's really just stupid," Scheiner says of the workstation's ability. "No matter how much processing I do, I just don't seem to run out of room." Scheiner pointed out that "AMD doesn't pay me." He's just grateful the technology has allowed him to build a professional-level studio in his home so he can spend more time with his family after a lot of years on the road working. His first in-home project was mixing an Aaron Neville album.

For the Dallas concert, Scheiner and Cherney recorded the audio in live, two-track stereo using four DAW64s housed in two very impressive mobile (and very air-conditioned) recording studios parked outside the steamy Cotton Bowl. The entire show was recorded to one drive, and as a backup, a recording was made to tape, but "that loses resolution," Scheiner says. The show was also broadcast by SIRIUS Satellite Radio and streamed across the Internet on a 15-second delay.

"We're going into this with the idea that everything is going to work fine," Scheiner told us before the show. The tricky part comes when it is time for the surround mixes, he said.

Clapton plays to a packed Cotton Bowl.

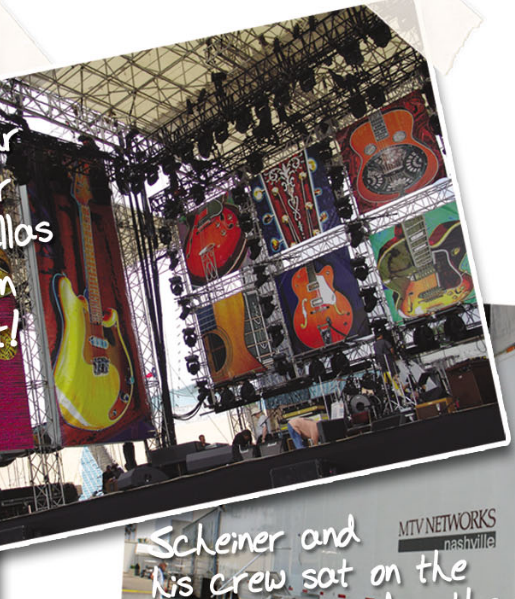

An 11-hour gig under the Dallas sun can get hot!

Scheiner and his crew sat on the sidelines recording the gig.

AMD Dual Opterons For The Digital Audio Generation

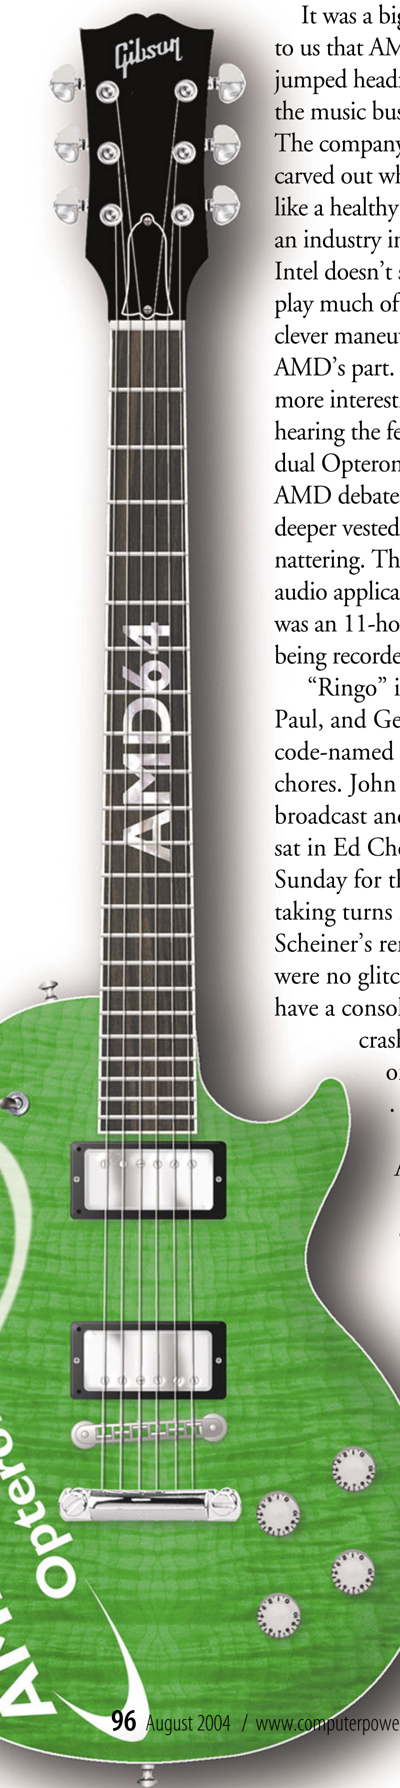

We sat with Scheiner for an hour talking technology and rock 'n roll.

It was a big surprise to us that AMD had jumped headfirst into the music business. The company has carved out what seems like a healthy niche in an industry in which Intel doesn't seem to play much of a part—clever maneuvering on AMD's part. Even more interesting is hearing the feedback from audio engineers about using DAW64s, which are dual Opteron 64-bit digital audio workstations. We've all heard the Intel-AMD debates between fanboys on message boards, but audio engineers have a deeper vested interest in the technology beyond the mine's-bigger-than-yours nattering. They need a technology that can go the distance for mission-critical audio applications. For example, the Crossroads Guitar Festival's finale show was an 11-hour event being mixed and broadcast live while simultaneously being recorded for future presentations. It was all being done on DAW64s.

"Ringo" is the DAW64 pictured on the previous page. However, John, Paul, and George aren't far behind. Together these four DAW64 boxes, code-named after each member of The Beatles, took on the recording chores. John took on duties Friday night to record a SIRIUS Satellite Radio broadcast and took care of the Saturday recording in the Dallas heat. Paul sat in Ed Cherney's remote trailer (the MTV one shown in the picture) on Sunday for the big show, while George hung out backstage on Saturday taking turns recording live performances with John. Ringo sat in Elliot Scheiner's remote trailer on Sunday taking turns with Paul. "There really were no glitches. Everything went pretty smoothly," Scheiner says. "We did have a console crash during one performance. Actually half the console crashed, so we got the entire band but missed the vocal. That was on Vince Gill. He ended up resinging, though, so it worked out . . . It was all so great, and it really sounds great."

We have already mentioned that DAW64s feature dual AMD Opteron processors. The other goodies include up to 16GB DDR, four hard drives with optional dual removable drive bays, and NVIDIA Quadro 4 380XGL for graphics.

Talking Music

When you're an audio/music fanatic (and we are), it's hard not to get caught up listening to someone like Elliot Scheiner's stories. Here are a few of our favorites:

♩ Fleetwood Mac's Stevie Nicks wanted to move into Scheiner's house for a few weeks to help him mix an album. Scheiner told her no. "Stevie travels with an entourage."

♩ Beck's "Sea Change" was recorded so well and was so beautiful that "it really didn't need much done to it."

♩ An initially pessimistic Walter Becker of Steely Dan was so impressed with Scheiner's mixes of the band's "Pretzel Logic" using his home-studio DAW64, Becker said, "I want one of these." He later got one for his home in Hawaii.

♩ The Honda Acura TL "was a real love thing," Scheiner says. "No format has succeeded until you got it in the car."

♩ Once in a music store, an employee demoed a 5.1 system for Scheiner using a Steely Dan album that Scheiner had mixed in 5.1, cautioning him first that "the quality isn't the best."

♩ Ray Charles is Scheiner's all-time favorite artist, but he never worked with him.

♩ Scheiner worked on Billy Joel's "Songs In The Attic," which involved the first use of digital live recording. "I thought CD [quality] was terrible. It didn't have the warmth," he says. "DVD-Audio is the closest thing to analog there is."

♩ Scheiner has won five Grammys. This year, a Grammy for best surround-sound recording will be awarded. Everything to date recorded in surround sound is eligible, making Scheiner a definite favorite.

Bring It On Home

As of press time, Scheiner was holed up in the studio putting in long hours. Warner Brothers will have its hands full choosing material, as Sunday's concert, as well as an all-day Guitar Village show Saturday on the Fairgrounds, were beyond description. (B.B. King, Buddy Guy, John Mayer, and Clapton trading licks for an hour? Come on.) As for AMD, the company is hoping the DAW64's abilities will eventually trickle down from the professional level to the garage band arena, giving up-and-comers a tool to record and mix their own songs with professional-level quality.

by Blaine A. Flamig & Samit G. Choudhuri

SOFTWARE TIPS & PROJECTS

Editing Audio Archives, Part 2

LAST MONTH WE USED SOME COMMONLY AVAILABLE SHAREWARE AND SOFTWARE BUNDLED WITH CREATIVE LABS SOUND BLASTER AUDIGY CARDS TO CLEAN UP digitized versions of old analog recordings. This month we look at some more advanced noise-reduction tools now available in a higher-end commercial product, namely the newly released Adobe Audition 1.5.

Veteran digital audiophiles may recall that Adobe purchased the old CoolEdit program years ago and turned it into Audition. Like a number of recent editing programs, Audition 1.5, which should be available as you read this, helps automate standard audio cleanup procedures, but it also offers a host of other more advanced techniques for bringing older and damaged audio tracks up to contemporary digital specs. While the specific instructions here apply to Adobe Audition, you can apply many of the principles in other programs.

Create Your Own Noise Filter

One of the most effective ways to eliminate repetitive background hum or buzzes on a track is to create a customized noise profile from that noise and then apply it as a filter to the affected area. Open the program and clear the screen of all other projects by choosing Close All from the File menu, saving any work you may want to preserve. Open the audio file you need to clean and make sure you are in Waveform view.

To make your own noise filter, you need to isolate about a second of audio on the track in which there is only background noise. Usually you will find this at the lead-in of the track, before music or dialogue begins. On old audio voice recordings, you should be able to find loads of pauses in the voices where there is a second or so filled only with the sort of background noise you want to eliminate throughout the track. Use the Zoom key (=) to help isolate that second or so of audio and make it large enough on the screen to highlight.

Click and drag your cursor to highlight the area. Using the Play button with an area highlighted will only play back that blocked section, so you can check to make sure that only background noise has been blocked. Choose Noise Reduction from the Effect menu and then click the Noise Reduction tool. In this window, click Capture Profile. This copies the noise you selected and essentially turns it into a filter that will pull that same noise out of part or all of your track.

To apply the filter to the whole track, click Select Entire File. The Noise Reduction Level is set to 100% by default, but if you find that applying the filter muffles or distorts the overall sound, you can try lowering the filter level. Try setting the slider to different levels and then use the Preview button to hear the effect it will have on the entire track. Once you have your desired settings, click OK to apply the filter.

WinXP Tip Of The Month

Retrieve Your Lost Briefcase

Windows XP users may well wonder whatever happened to that My Briefcase file-synching utility from previous Windows versions. In fact, it is still here, albeit less obvious than it once was. To put the familiar My Briefcase icon on your Desktop, simply type **syncapp** in the Run dialog box. Alternatively, you can right-click the Desktop and choose the New and Briefcase commands to make a new Briefcase.

The first time you open one of these folders, Windows will outline how to use it to synchronize files between two PCs. Basically, you drag files into this folder and move the Briefcase to a floppy or across a network for use on another machine. After revising and saving the files on another PC, put the Briefcase back on the original machine and open the folder. You can click one or more of the files and tell Briefcase to apply these changes to the original files. ▲

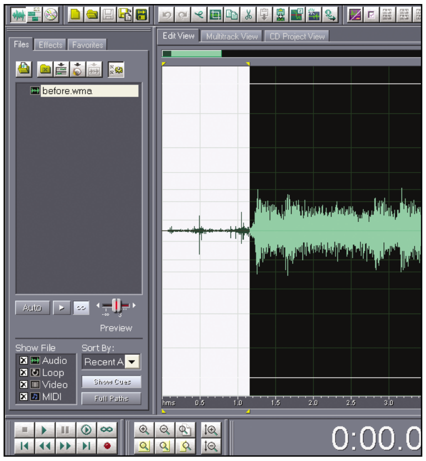

Audition lets you turn a passage of pure background noise from a song lead into a filter that removes the noise patterns from the entire track with minimal effect on the music itself.

To apply the custom noise filter only to a section of a track, close this window. In Waveform view, locate and highlight the area to which you want to apply the noise reduction, and then reopen the Noise Reduction window. Click OK to apply the filter only to the highlighted section. Remember that you can use the Undo command in the Edit menu to back out of this filter. If you don't like the results of your filter, start again from scratch by using the Undo command.

As we mentioned in the first part of this series, it is always best to apply one audio cleanup fix at a time to a file, so once you apply a custom filter to your liking, save the file using a new file name. Use this new file as your future baseline from which to apply the next treatment, such as click/pop removal. This approach also provides you with a set of backups in various states of refinement if you decide later to redo a cleanup project from a particular stage.

Say It Loud

Before video camcorders were the norm, many people used old cassette and reel-to-reel tape recorders to preserve events and the voices of family members. But poor recording equipment and tape degradation often leave the voices on these recordings muffled and indistinct. Digital filtering and adjustments of certain frequency bands can help poorly recorded voices stand out from background noise and even make previously inaudible passages decipherable.

In Adobe Audition, open the track you want to fix. Use the Effects menu to access Filters and Graphic Equalizer. Click the 10 Bands (1 octave) tab at the top of the Graphic Equalizer dialog box so that you are only dealing with 10 frequency sliders. Most other audio editors also offer eight or 10 band equalizers to apply to your tracks, so you can use the basic principles here outside of Adobe Audition. In most recordings, the voices will lie in the 1KHz (1,000Hz) and 2KHz (2,000Hz) frequency ranges, but adjusting the other bands can also depress background sounds to help the voices poke through.

Use the Preview button to play the track and experiment with equalizer settings on the fly. In old voice recordings, we start with all bands slid to the bottom so that the track is barely audible. Experiment with the sliders to locate the bands containing most of the vocals (usually 1KHz and 2KHz) and bring them about three-fourths of the way up. Remember that moving the bands too far up can introduce distortion, so don't be overzealous. Once you have decent vocal volumes, bring up the surrounding bands only so far as they fill out the vocals without adding extraneous noise. When you have found your best settings, click OK button to apply the new equalizer settings to the entire track. Now you can save the track in whatever format you like.

Erase A Jerk

A lot of concert audio contains irritating audience noise, such as sickly listeners with insistent coughs or the one jerk in the crowd who wants to make sure he gets on the band's "live" album by whistling or shouting during quiet passages. With the most advanced tools now available to consumers, you can actually isolate the frequency band that contains these rude intrusions and eliminate them without harming the rest of the soundtrack. Audition calls it "frequency space editing."

Begin by locating the offending noise in your audio track. Now use the View menu to switch into Spectral View. In most cases the intrusive noise will look like a smudge that towers over the rest of the track. If you are having trouble seeing this audio smudge, open the Settings screen (F4) and click the Display tab. Use the Resolution drop-down menu to raise the resolution a couple of steps. Click OK to make the change.

Once you have located the noise on the spectrum, use the Marquee (M) tool to highlight it. Now go to the Favorites menu to choose Repair Transient. This will erase almost all of the "transient" sound and leave the rest of the track unaffected. CPU

by Steve Smith

Registry Tweak

Open In Notepad

Whenever you come across a poorly named file that has no program association or identifying file name suffix, your best bet is to try to open it with Windows XP's built-in text editor Notepad. This Registry tweak lets you right-click any file and access an Open With Notepad command instead of drilling through the usual nest of Open With submenus. In the Registry, go to HKEY_CLASSES_ROOT*, right-click and select New and then Key, and name the new key Shell. On some configurations, this key may already exist. Now, create another key under Shell and name it Open. Double-click the Default value for Open in the right pane and type **Open with Notepad** as its value. Create another key named Command under the Open subkey. Double-click the Command key's Default value and type **notepad.exe %1** as its value. The effect should be immediate; you should now be able to right-click a file and see an Open With Notepad command. ▲

Infinite Loop

Missing Gene Means Bad Hair Days

Ever feel like your hair is out of your control? It may be, literally. Researchers at Johns Hopkins School of Medicine and the Howard Hughes Medical Institute have found that the Frizzled 6 gene determines hair patterns. By removing the gene in mice, scientists concluded that its absence causes unusual hair whorls and cowlicks. A nearly identical gene appears (or doesn't appear) in humans, suggesting that genetics play a role in unruly hair patterns. ▲

Source: www.cnn.com

WARM UP TO PENGUINS

Monitor Your Hard Drive Space

AS HARD DRIVES CONTINUE TO GROW TO CAPACITIES WE CAN'T EVEN COMPREHEND FILLING, WE SOMEHOW STILL MANAGE TO. RATHER THAN wait until you fill that space, however, automate your system to warn you when space is getting too close for comfort. You can then farm off those old digital home movies that are taking up a huge chunk of your drive to CDs or DVDs to free up room. Before we detail how to do this, we need to look at the commands that help track down drive-usage problems. Then we'll build a script you can use on your own system for automation.

How Much Space Do You Have?

The primary command for tracking available and used hard drive space is df, though most people use it in the form df -h. The -h portion tells the command to report the space in values of gigabytes or megabytes instead of kilobytes. In Figure 1 that accompanies this article, there are two hard drives listed, hde and hdf. The output columns from left to right are:

- Filesystem—the partition the line refers to
- Size—how large the partition is
- Used—the space that is used
- Avail—the space free on the partition
- Use%—the percentage of the partition currently full
- Mounted on—the partition's mount point, which is where the partition is on your directory structure

As Figure 1 shows, the system has a root partition mounted on /dev/hde2, a /boot partition mounted on /dev/hde1, and something called /mnt/FC1 mounted on /dev/hdf3. /mnt/FC1 is a temporary mount containing old files from a Fedora Core 1 installation, so we're not interested in how full or empty that particular partition is. The /dev/shm listing with the partition "none" is interesting, however. Shm stands for shared memory, in which a program stores data in RAM, allowing other programs to directly access that data from your machine's memory rather than from the drive. This is great for fast data access, but it has nothing to do with keeping our drives from overflowing, so we can ignore the /dev/shm line.

This leaves us with two lines:

```
/dev/hde2  54G 5.3G  46G 11%  /
/dev/hde1  99M 5.9M  88M 7%  /boot
```

Neither of these partitions is even half full, so there's nothing to worry about for now. As you update your system over time, the /boot partition will slowly grow, especially if you don't remove old kernels. The root partition can fill much faster, especially if you download or make lots of multimedia files and keep them.

Building A Warning Script

A script is a task that's simple enough we don't need an entire shell script. Instead, we just need to figure out the command we want to use and plug it into our system's automation software by itself.

To begin, we know we want to run the du -h command. However, the purpose of this script is to let the computer monitor how much space we have left, only bothering us when the partition(s) begins to fill up. We need to decide what the danger point will be. Let's say when a drive reaches 80% full, we want to clean it up.

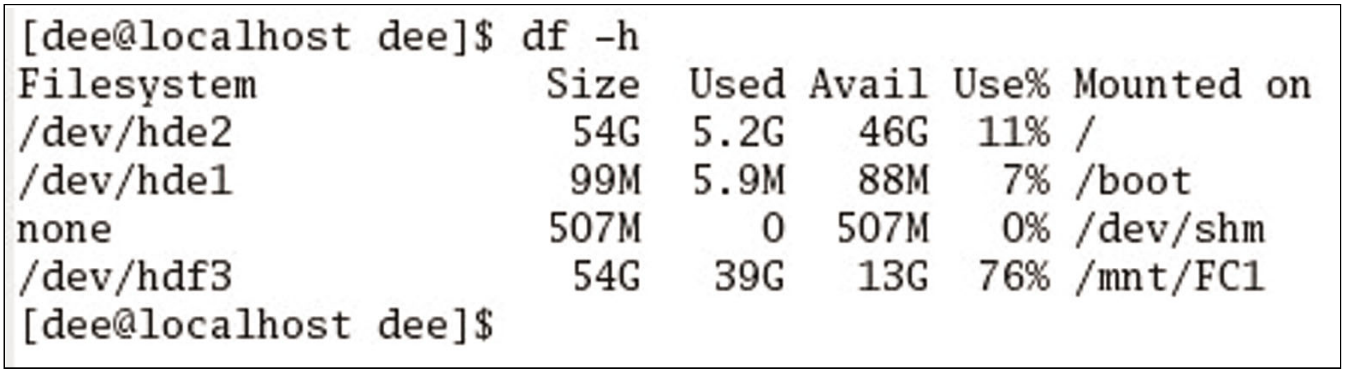

```
[dee@localhost dee]$ df -h
Filesystem            Size  Used Avail Use% Mounted on
/dev/hde2              54G  5.2G   46G  11% /
/dev/hde1              99M  5.9M   88M   7% /boot
none                  507M     0  507M   0% /dev/shm
/dev/hdf3              54G   39G   13G  76% /mnt/FC1
[dee@localhost dee]$
```

Figure 1: This illustration shows sample df -h output for a Fedore Core 1 installation. Two hard drives, hde and hdf, are shown.

There are many ways to approach building this script. Some scripts are incredibly complex. Our purposes are more practical. We just want a warning and no more, so we'll make things simple. For starters, we want to use the df -h combination. Then we can take the output from this command and, in the Unix way, use a pipe (|) to feed the output as the input to another command, namely grep. The grep command searches for particular text, such as “80%.” So, our script might look like this:

```
#!/bin/bash
df -h | grep 80%
```

There's the potential to make a major mistake here and completely miss a big spike in drive use, however. What we want is any value that's 80% or above. The question is how to tell the grep command to search for anything from 80 to 100 with a % at the end. We need to use regular expressions, which we discussed in June's issue (page 99). It's actually more complicated to look for three digits instead of simply two, so we're going to ignore the 100% mark. Instead, we'll look for usage between 80% to 99%.

How do we do this? We know that we're only interested in a first digit of 8 or 9. You can specify 8 or 9 in a regular expression by using [89]. After the 8 or 9 is a second digit (making the number 81 or 91 up to 89 or 99). It's important that we make this a second digit only and not a third, fourth, etc. Because the second digit can be any number from 0 to 9, we can specify this as [0123456789]. Now we need to add the %, so our regular expression looks like:

```
[89][0123456789]%
```

That means our command is:

```
df -h | grep [89][0123456789]%
```

Automating The Space Test

To be useful to us the command needs to run at least once a day. We can set this by adding the command to a cron job. Cron is a daemon that runs regularly scheduled tasks on most Linux machines, regardless of the distro you're using. You don't need to be root to run this, so you can create a regular user cron job by typing **crontab -e** at a command prompt. This opens an empty file where you can store as many cron jobs as you like. Your default text editor will govern how you can edit crontab files. Most likely, the default text editor will be vi or emacs.

The first portion of a cron job is the time(s) the program should run. A crontab file has a flexible set of possible timings. We can't cover them all here, but you can type **man 5 crontab** at a command prompt to read the options. We'll stick with the basics for running a job once a day, say at 12 a.m. A crontab entry looks like this:

```
MIN HR DM MON DW command
```

MIN is the minute, HR the hour in 24-hour time, DM the day of the month, MON the month, and DW the day of the week. To indicate “all,” you will use an asterisk. So, to run the same command every day at midnight, the MIN is 0, the HR is 0, the DM is “all,” the MON is “all,” and the DW is “all.” This will give us the following (in crontab-speak):

```
0 0 * * *
```

Now we plug the command at the end:

```
0 0 * * *        df -h | grep [89][0123456789]%
```

Once we save and exit the file, we have an early warning system. This command works properly at the command prompt, but it won't work quite right from cron due to various shell expansion rules, as the characters process a bit differently from cron. To fix this, change the command to:

```
grep -e “[89][0123456789]\%”
```

The -e flag for grep tells the system to look for regular expressions, and putting the whole value in quotes tells the system to take everything literally. Adding the \ in front of the % makes sure the % isn't interpreted incorrectly, which is what will happen if you don't make this change.

When the command runs, any output is emailed to your shell account. Note that output won't go to your main email account unless you instruct this to occur. To do this, expand the line so that it looks something like:

```
0 0 * * *        df -h | grep -e “[89][0123456789]\%” | mail -s “WARNING!” you@example.com
```

Don't press RETURN if your code wraps down to the next line or cron will think it's looking at a new cron job. The mail command is a basic method for sending email in Linux, and the -s flag tells the command to set a subject line using the item in quotes. After that, replace you@example.com with your own email address.

This Is A Warning

There's little worse than suddenly discovering you have little drive space left. Protect yourself from having to boot with your system rescue media and clean out the drive under duress by automating a handy system babysitter. CPU

by Dee-Ann LeBlanc

Infinite Loop

Pringles Trivia

Munching on junk food is a typically mindless experience. That's about to change. Proctor & Gamble has invested in a new technology that uses food coloring to print trivia questions on Pringles snacks. An inkjet-like printer with a conveyor belt prints questions on one side and answers on the other side. For now, Pringles will be adorned with blue or red ink, with other colors to be introduced later. Who knew that eating Pringles could make you smarter?

Sources: ajc.com

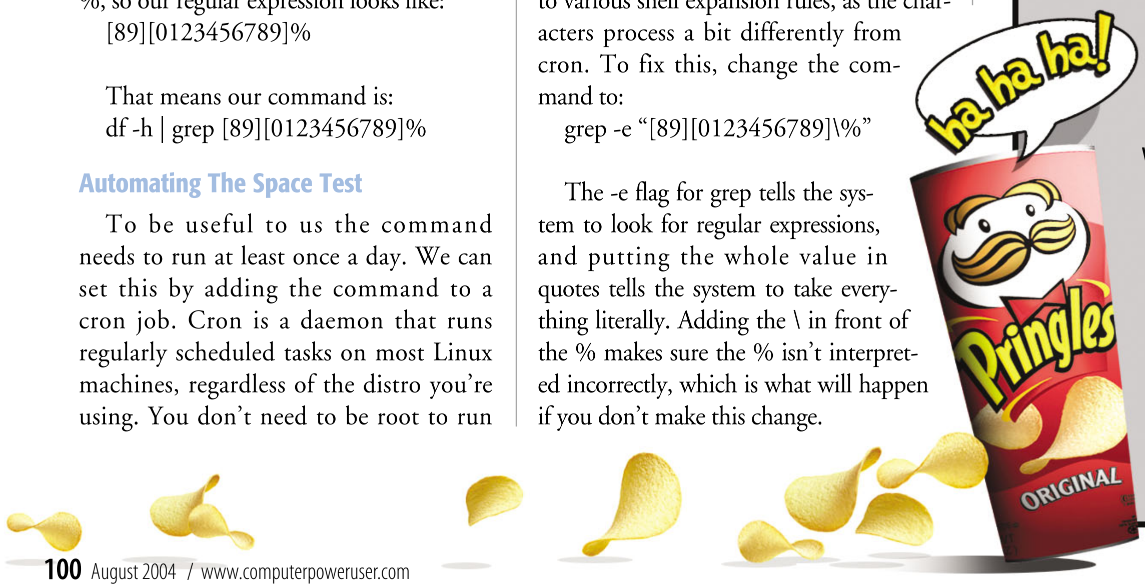

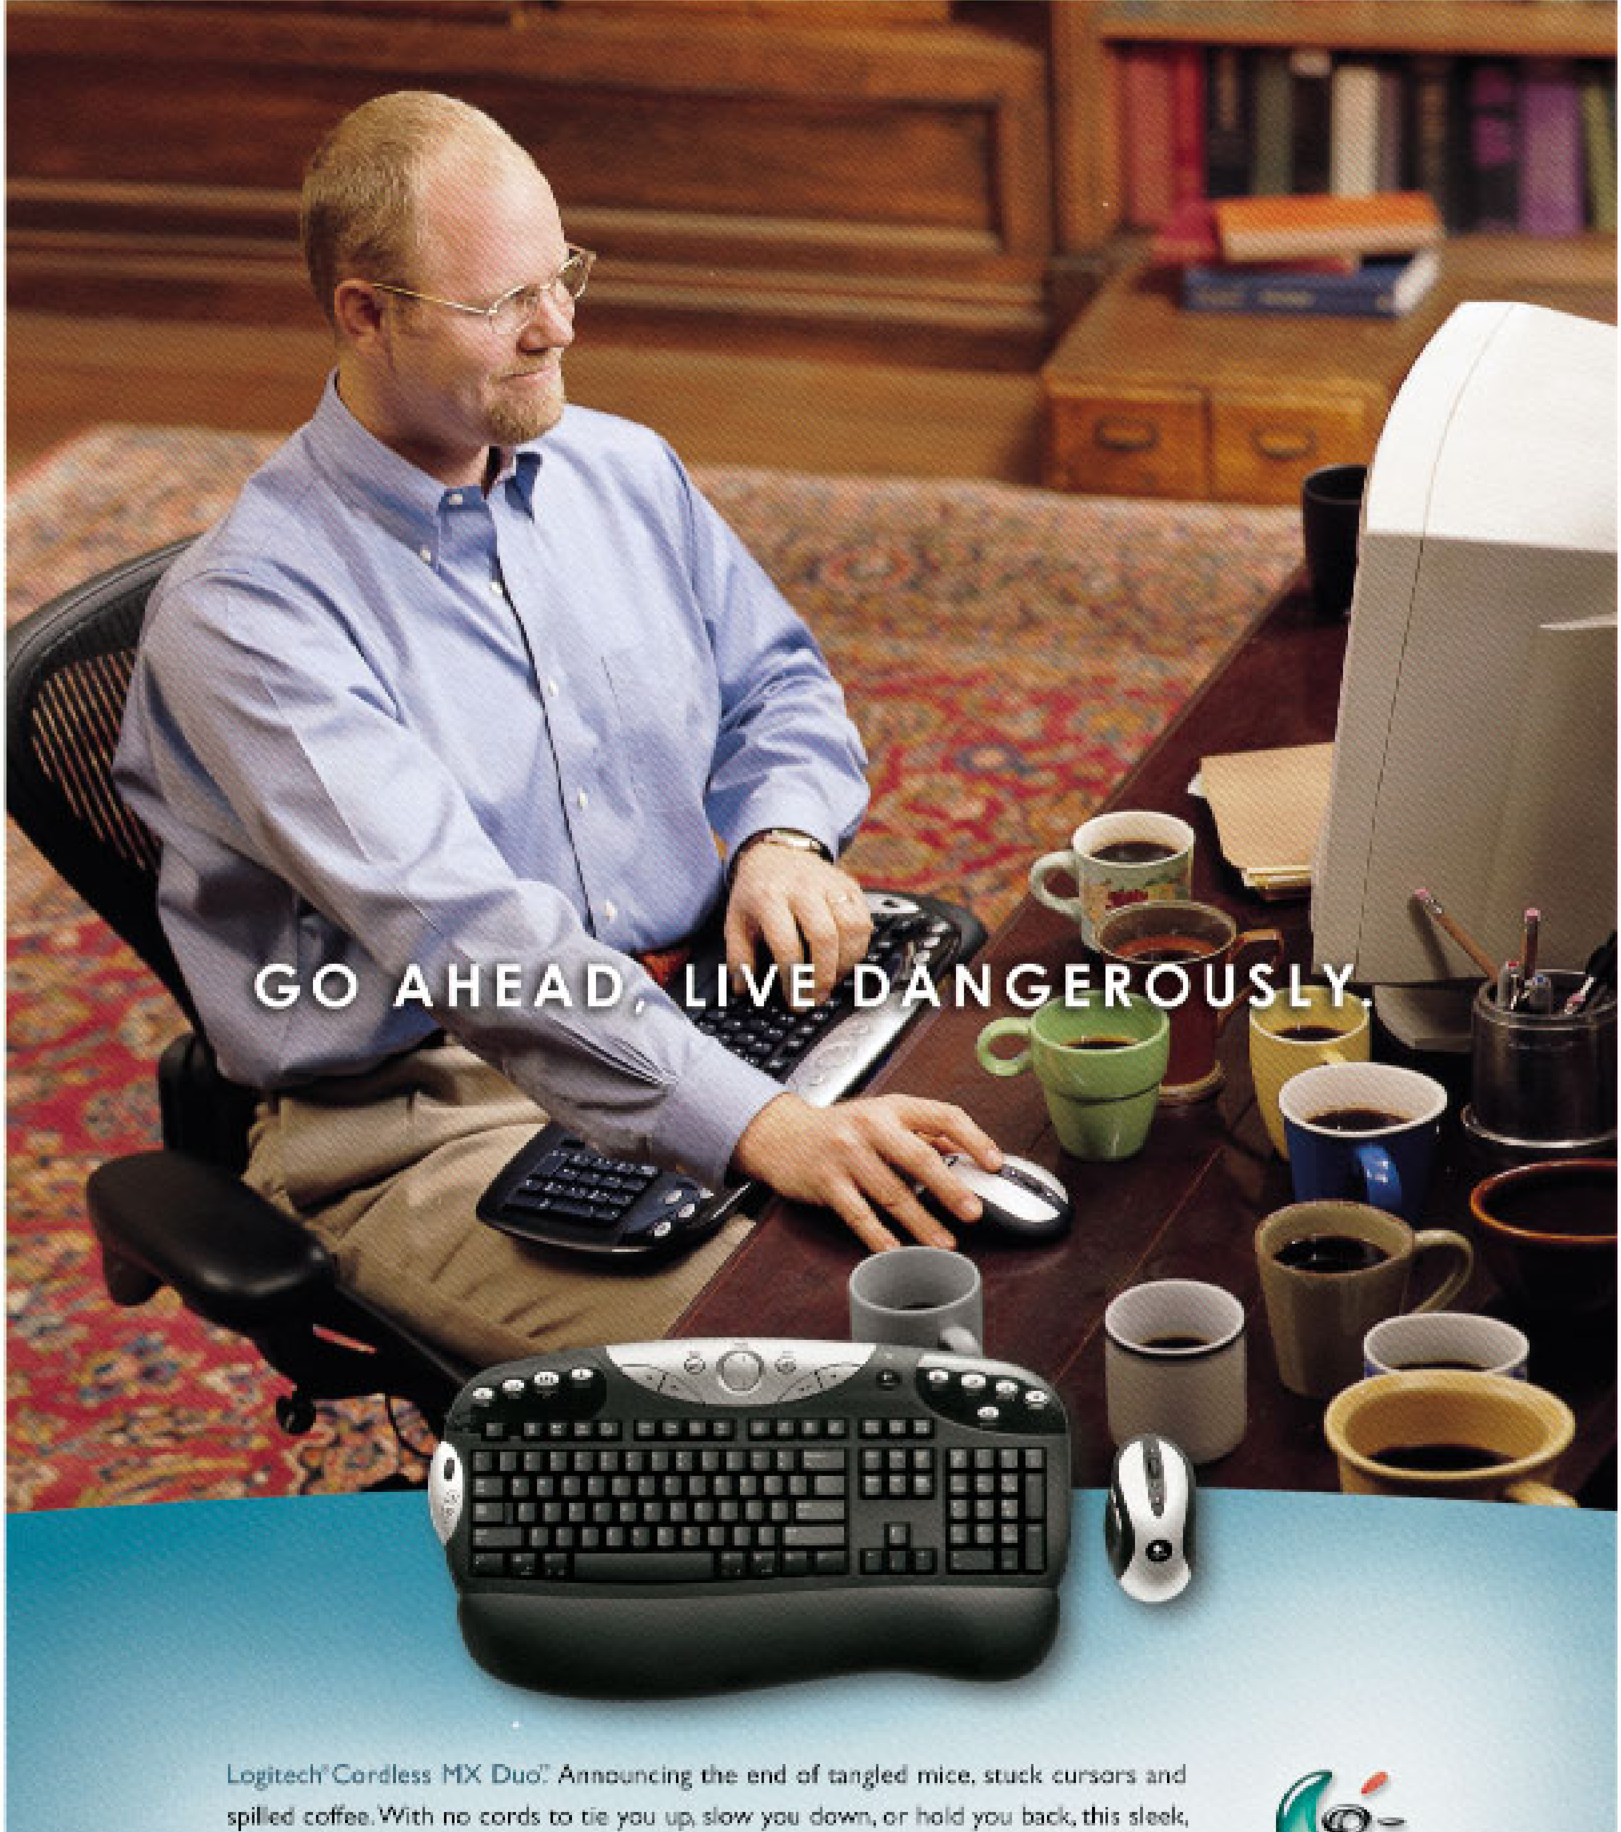

Logitech® Cordless MX Duo™. Announcing the end of tangled mice, stuck cursors and spilled coffee. With no cords to tie you up, slow you down, or hold you back, this sleek, one-touch keyboard and advanced optical mouse give you the control and flexibility to work your way – whatever the obstacles. What will you do with all that freedom?

www.logitech.com

Logitech

©2003 Logitech

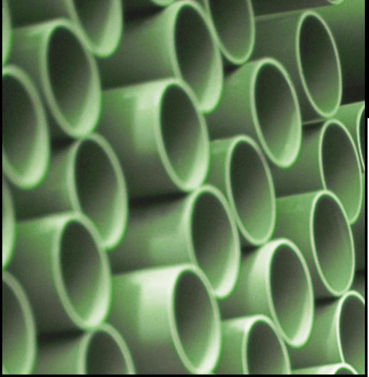

Shavings From The Rumour Mill

by Mike Magee

Will Multicores Save The Game For Intel & AMD?

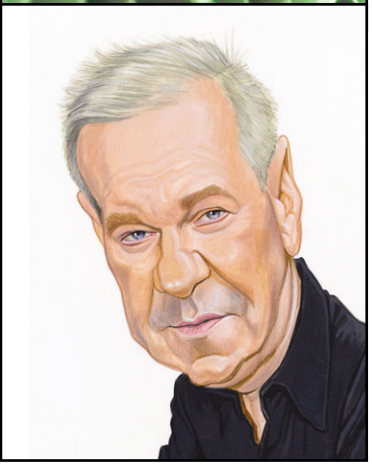

Mike Magee is an industry veteran. He cut his teeth on ancient products like the Dragon and the Japanese PC platforms long before the IBM-PC won. He worked for a corporate reseller in the mid-'80s and saw the Compaq 386 sandwich box and every GUI known to humankind. Mike decided that the way to go was the Interweb around 1994 after editing PC mags in the late '80s and '90s. A co-founder of The Register, Mike started the chip-driven INQUIRER (www.theinquirer.net) in 2001. He has contacts from top to bottom in the business, spanning the entire chain, who help him root out interesting rumours and speculation.

Santa's little workshop is, we are told, populated by legions of gnomes who spend night and day with their little hammers and chisels knocking out the toys to make the good children happy. Whenever we visit Taiwan, we look at the people who work hard night and day to knock out IT products, but unlike the benevolent figure of Santa, instead there's a whole team of ruthless Intel marketing men who spend their time going around trashing all the Taiwanese workers' efforts.

If we have two cores on one chip, does this mean that Microsoft will want to charge us two software licenses?

This was brought home to us in June when we visited the booths of such hardware luminaries as Asustek, Gigabyte, Abit, and ECS. It's not just Intel this year which is trashing the Taiwanese firms' toys. ASRrock is obviously so desperate about changes AMD is making to its chips that it even showed off a board for 754 and 939 pins. The legless Socket T for Intel 775 Pentium Prescott processor has also caused the Taiwanese makers a lot of grief. Intel told them they couldn't include AGP where there was PCI Express, that it was impossible to mix up DDR and DDR2 memory, and other stuff, too. We imagine the intrepid Taiwanese engineers scratching their heads for a few seconds and then saying, "Oh to heck with it; we'll do it anyway."

No sooner have the mobo companies spent a huge amount of money developing LGA775 and Grantsdale chipsets than Intel goes off and broadcasts to the world that performance doesn't matter now, and everything will be multicore next year and for a good time after. Most of the top chip pundits agree that the X86 architecture will soon reach a performance ceiling, although it appears that Intel has reached it first. As early as last year, Intel prepared us for this by positioning its Centrino Pentium M as a chip without any megahertz attached. Now, and more recently, the Pentium Prescotts have numbers, too, but they're not raw megahertz ones and bear a relationship to each other. Intel says that the new numbers represent features of a chip rather than overall performance.

The days of megahertz madness are now over, and we're moving into the era of multicores. The first company to introduce these is likely to be AMD rather than Intel. If we have two cores on one chip, does this mean that Microsoft will want to charge us two software licenses? A trickier question is that as the old rules change to the new, will this mean that it's the end-game for squeezing every extra ounce of performance out of your gaming or enthusiast machine? We all know, and have known for many years, that the hardware has pushed so far ahead of the software that it's only die-hard gamers and those editing videos that are going to need that raw computing power all at once. The chip companies, when they're pressed, can only come up with these two applications to justify a home user buying a 64-bit system with gigabytes of memory and a graphics card at a price that could get you a reasonable holiday in the Mediterranean or the Caribbean.

Dirk Meyer, who fathered the Athlon 64 architecture, got flustered when pressed on this point at Computex. He really racked his brains but couldn't come up with any applications that would excite us all. Then there's Longhorn, Microsoft's up and coming operating system. That was supposed to be out by now and was supposed to demand all this PCI Express, DDR-2 stuff that is rapidly becoming available. When it does come out, we are still promised a miracle OS which, in Bill Gates' words last year, costs more to develop than the initial Man on the Moon mission.

So we're all going to spend this money on hardware that does everything just to run an operating system from Microsoft? ■

Send rumours to "Mad Mike" Magee at Mike@cpumag.com.

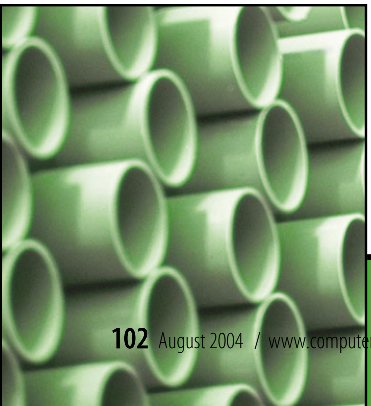

Technically Speaking

An Interview With Charlie Boswell, Director of AMD's Digital Media & Entertainment Group

If you can play "Chopsticks" or power chord your way through "Louie, Louie," chances are you've fantasized about being a "serious musician." Personal computers and MIDI have enabled many of us to fiddle with those fantasies, but eventually you're bound to hit a technical brick wall. There's only so much you can do on a desktop system—or was. AMD believes it can start musicians of all skill levels humming a new tune with its AMD64 platform. We sat down with Charlie Boswell, director of the digital media and entertainment group at AMD, to learn why ▲

by William Van Winkle

AMD's Charlie Boswell (right) surrounded by technology in the studio–and Dweezil Zappa.

CPU: What's the relationship between CPUs and audio?

Boswell: Just like the external modems from a decade ago can now be run in the CPU as software, today's host-based signal processing is occurring on audio faster than anything else. Some examples: virtual instruments, virtual synthesizers, plug-ins for digital audio in these digital audio workstations. We use the term 'digital audio workstation,' or DAW, to describe a PC or workstation that uses a digital audio application like Nuendo from Pinnacle Steinberg to collect samples of digital audio through some external hardware. You take those samples and string them to the hard drive, doing host-based signal processing on that audio. Consider it a digital tape machine but with lots of intelligence–all kinds of special effects, normal processing functions like compression, reverb, equalization, limiting.

In order for this to be successful, the platform must be very robust and able to handle lots of these pieces of software. Take what we did recently at the Crossroads Festival. We were tracking in 48 channels, or streams, of audio, each recording at DVD-Audio levels. Your current CD uses 16 bits of resolution per sample that was sampled at 44.1KHz. DVD-Audio is 24 bits sampled at 96KHz, so it's oversampled, yielding more of the harmonic nuance or sonic integrity of the audio. The software must be able to handle the dynamic range of the arithmetic in order to preserve sonic integrity.

The dual Opteron, the AMD64 Direct Connect technology, is being very successful here for a number of reasons, not the least of which is it provides for the fact that memory and I/O accesses don't compete for bandwidth. They're on separate wires. Also, low latency is incredibly important in this application. In a sample data system, you can't drop or skew samples. Our Direct Connect technology has the lowest system latency available today. It has no frontside bus. Each processor directly communicates with memory. And each processor directly communicates with the other processor.

CPU: That's the Opteron white paper. What does it mean in English?

Boswell: It means that you wind up with balanced system architecture. In an audio application, as I begin to load up tracks and effects and plug-ins, I'm loading up the processor. If I have a processor that performs statistically and has a frontside bus like the dual Xeon, I'm not going to be able to do this.

CPU: What do you mean by 'statistically?'

Boswell: It's another word for unpredictable behavior. Ethernet has a statistic behavior. If you have collisions on Ethernet, you have what is called an 'exponential back-off.' It is like two people trying to talk at the same time. One person has to back off and let the other talk. On the Intel frontside bus, data are fighting for bandwidth, so I wind up having some unpredictable behavior. In a real-time situation like audio, this is a killer.

But with an Opteron, I can load this thing up to 90+% and not suffer any dropouts on the audio. This is incredibly important for these types of real-time media-creation applications. This allows me not to have to take any intermediate steps. Like, say I determine that my processor wouldn't allow me to have all these tracks and effects. I might have to do some intermediate rendering of those tracks, making a sort of work-around artistic judgment, if you will, that might not be in the best interests of my final creation. But if I'm allowed to have everything in real-time on the digital audio workstation, I don't have to make any mix-down decisions that affect my product.

CPU: Video can be just as much if not more of a data-heavy real-time app as audio, yet Intel often claims superiority in video apps. Is that just marketing?

Boswell: Audio is a lot less forgiving than video. We have this thing in our head called 'persistence of vision.' You can drop frames in video, and it's not as discernable as audio. If I've got the wrong sample in audio, you're going to hear a click or a pop. In video, you have enough persistence of vision not to perceive that stuff. Now, Opteron has the same robustness in audio and video. We can load our system up much higher than those guys can without any kind of break-up.

CPU: That's all well and good if you're a sound engineer. But from our perspective as a music listener, what difference does this make to our experience?

Boswell: Well, the guy who's creating the surround audio has the ability on an AMD to add more effects, more tracks. He can go farther than with the platforms he had before. Likewise, that same architecture can be used by a kid for his garage band. It's essentially the same technology. He can record in the home and use more virtual instruments. There's a company called Open Labs where you'll see part of that revolution in a product called the neKo 64, which is a workstation in an instrument. Guys like Prince and Dr. Dre and The Who are using this on tour right now.

CPU: What is a neKo 64?

Boswell: Imagine a workstation PC with a synthesizer keyboard and a flat-panel touchscreen built into its case running these host-based audio programs. Knobs on it are used to control different filters, different voices. The keys send note-on, note-off commands to the system. You can do scratching on the touchscreen, which you can't do on a laptop. It can also record with Nuendo software while you're performing. The processor is behaving as the synthesizer. Musical instruments have adopted an open architecture, the PC. But you can't just run those instruments on an Intel box because you're adding lots of effects and expecting to load that processor without it failing. The Xeon can't do that. Hyper-Threading can't do that.

CPU: So can Prince, out there on stage, do things with this instrument that he couldn't do before? Is 'Little Red Corvette' in 2004 going to sound different than it did in 1984?

Boswell: Look at it in terms of performance. A synthesizer has a fixed set of voices that it can do. A sampler would be the next way to go—you load in new samples. Or you put an open architecture into a musical instrument, and now he's unlimited in what he can do. He's got programs that can generate new synthesis. He's got unlimited voices. As a performer, he can realize anything sonically. Companies like Native Instruments let you download new patches and voices, and you can create new things on-the-fly. But for all that to work, you have to have an architecture that's not going to freeze up and fall on its butt in front of the crowd.

CPU: I see on Open Labs' page it talks about 64-bit processing, but all of your commercial DAW systems today are running under 32-bit Windows. So how much of a factor—if at all—is 64-bit computing in digital audio?

Boswell: It's huge, especially when you consider the apps these people are running. I'm thinking of apps like GigaSampler. In a 32-bit process, you have a maximum of 2GB that process can have. The operating system gets the other 2GB. When you go beyond that, you start swapping to rotating media, like the hard drive. That's bad for real-time things. So now when I open up 64-bit addressable memory, I can have multiple samplers, multiple virtual instruments, with tremendous production capability at my fingertips.

CPU: Yet you need a 64-bit OS to get that additional memory access, and you're still running on Windows XP.

Boswell: Well, we have the beta on that OS, and we're doing a lot of testing with those instruments. And with Microsoft's OS coming later this year, lots of that software is already taking advantage of that extra memory.

CPU: What do you lose in stepping back to a single processor?

Boswell: It depends on what you're doing. If it's render time on movie creation, dual processors certainly scale very nicely. In audio, it's just the number of tracks and effects per track or the number of virtual instruments you can run.

CPU: If studio technology is now a home technology, what kind of impact do you think that will have on the overall quality of music going out to consumers?

Boswell: I think it's going to go through the roof. For starters, you're drafting in more talent. Ultimately, the public and the record companies decide who succeeds, but this is opening up an opportunity for new voices. It's providing the tools for young and old minds, people that are on the fence. People who want to express themselves can afford now to do it. They don't have to pay a studio $150 an hour. I'm sure that's squashed a lot of potentially great songs. CPU

To read our entire interview with Charlie Boswell, go to
www.cpumag.com/cpuaug08/boswell

What's Everyone Talking About?

ComputerPowerUser.com

Take Advantage Of All The Extra Benefits Of Being A Paid Subscriber Today!

Here's what's available to everyone:

- The CPU Web Log. Each weekday, our editors scour the Web to bring you the best high-tech (and occasionally offbeat) information in an entertaining way.
- Forums to share (or show off) your thoughts and expertise. Great whether you know it all or want to know more.

Subscriber Tip Of The Month

My Personal Library! New Features

Read an article you need to print out or refer to in the future? My Personal Library solves that problem! As a subscriber, you will see an Add To My Personal Library link in the upper right of ANY full article from *CPU* and our sister publications. Click the link to add the article to your library. You are now also able to customize your Library display features, by using collapsible and expandable folders. Click the folder to view all content within the folder, and click again to hide content. You will be able to permanently customize your Library display view by clicking the Set Library Options link in the upper right. Choose which setting you prefer to use and click the Change button. From that point on, your library display will always default to your customized option. Don't forget to check out this great new feature!

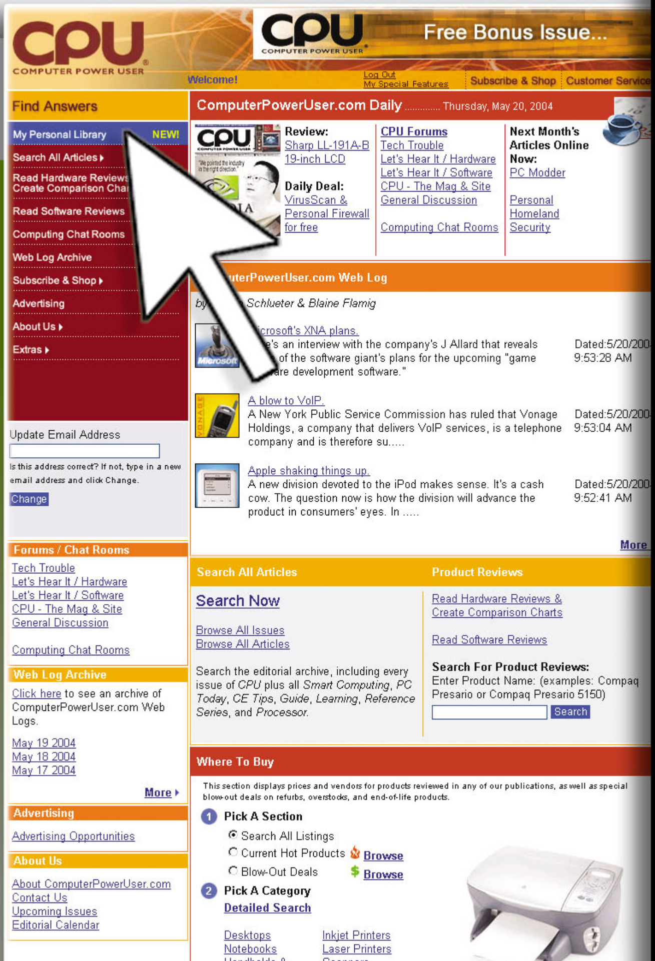

ComputerPowerUser.com Daily .... Thursday, May 20, 2004

Review: Sharp LL-191A-B 19-inch LCD

Daily Deal: VirusScan & Personal Firewall for free

CPU Forums
Tech Trouble
Let's Hear It / Hardware
Let's Hear It / Software
CPU - The Mag & Site
General Discussion
Computing Chat Rooms

Next Month's Articles Online Now:
PC Modder
Personal Homeland Security

uterPowerUser.com Web Log

by Schlueter & Blaine Flamig

icrosoft's XNA plans.
e's an interview with the company's J Allard that reveals of the software giant's plans for the upcoming "game are development software." Dated:5/20/200 9:53:28 AM

A blow to VoIP.
A New York Public Service Commission has ruled that Vonage Holdings, a company that delivers VoIP services, is a telephone company and is therefore su..... Dated:5/20/200 9:53:04 AM

Apple shaking things up.
A new division devoted to the iPod makes sense. It's a cash cow. The question now is how the division will advance the product in consumers' eyes. In ..... Dated:5/20/200 9:52:41 AM

More

Update Email Address
Is this address correct? If not, type in a new email address and click Change.
Change

Forums / Chat Rooms
Tech Trouble
Let's Hear It / Hardware
Let's Hear It / Software
CPU - The Mag & Site
General Discussion
Computing Chat Rooms

Web Log Archive
Click here to see an archive of ComputerPowerUser.com Web Logs.
May 19 2004
May 18 2004
May 17 2004
More

Advertising
Advertising Opportunities

About Us
About ComputerPowerUser.com
Contact Us
Upcoming Issues
Editorial Calendar

Search All Articles
Search Now
Browse All Issues
Browse All Articles

Search the editorial archive, including every issue of *CPU* plus all *Smart Computing*, *PC Today*, *CE Tips*, *Guide*, *Learning*, *Reference Series*, and *Processor*.

Product Reviews
Read Hardware Reviews & Create Comparison Charts
Read Software Reviews

Search For Product Reviews:
Enter Product Name: (examples: Compaq Presario or Compaq Presario 5150)
Search

Where To Buy

This section displays prices and vendors for products reviewed in any of our publications, as well as special blow-out deals on refurbs, overstocks, and end-of-life products.

1 Pick A Section
● Search All Listings
○ Current Hot Products Browse
○ Blow-Out Deals Browse

2 Pick A Category
Detailed Search

Desktops
Notebooks
Inkjet Printers
Laser Printers

Only paid subscribers to *CPU* magazine get complete access to all these great features:

Latest Issues. Get a guided tour through the latest editions of our technology publications.

More hands-on hardware reviews. Are the latest products worthy of the marketing hype? We put products to the test, complete with comparison charts and links pointing to where you can buy the products online.

Even more great reading. Our Web site features every article from *CPU* and a searchable archive of more than 18,000 articles from our sister publications *Smart Computing*, *Learning Series*, *Guide Series*, and *Reference Series*. It's like getting five magazines for the price of one!

A showcase of articles from next month's issue. Impress your friends with all the insider information you get by reading *CPU's* latest articles and product recommendations before they're available in print.

CPU COMPUTER POWER USER

Under Development

A Peek At What's Brewing In The Laboratory

IBM: Read My Lips

When you're trying to talk (shout) with somebody at a party or dance club, your eyes are usually glued to their lips. You might not be able to read lips per se, but the extra information you glean from watching that person's mouth movements helps to supplement the information your ears receive. When someone walks between the two of you, observe how your comprehension level plummets. Roberto Sicconi, manager of multimodal conversational systems, leads a team at IBM devoted to bridging the worlds of audio and video to develop improved speech recognition and reproduction systems.

"What we found is that lip reading alone does not give you the same level of accuracy that audio recognition does," says Sicconi, "but the features that you can extract are complementary to the ones that you get from the audio. For instance, some of the sounds you get in the audio domain look different in the visual domain. Ps and Bs and Ns and Ms have similar sounds that can confuse an audio recognizer, but they look completely different to the camera. By combining these two channels, we can restore a good deal of accuracy, even in very noisy conditions. We've proven with prototypes that we can gain the equivalent of 10dB in signal-to-noise ratio, so the noise can be 10 times as loud, and you still get the same level of accuracy if you have visual information accompanying the audio."

Sicconi and team currently has its prototype systems running on high-end notebooks, but IBM is looking at creating customized DSP chips for a much lower-cost implementation. This would enable not only cars to include the A/V speech system but also camera phones, PDAs, and other corporate products.

Once visual recognition becomes accurate enough, there may be situations in which you could dispense with speaking altogether. "You could whisper your PIN code or even just mouth the words to get the system to recognize it," says Sicconi.

According to Sicconi, his group's A/V technology will take another two to three years to perfect and then another two years to find its first market implementation in cars, although there could be some overlap between the time frames. ▲

IBM researcher Makis Potamianos works with the company's prototype audio/visual speech recognition system, which essentially lip reads to assist recognition accuracy.

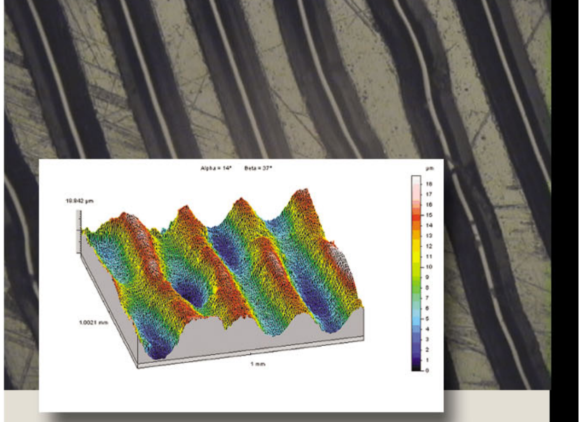

Lawrence Berkeley Tool Exhumes Old Recordings

Lawrence Berkeley National Laboratory physicists developed their special optical metrology tools for use in inspecting alignment in silicon detectors destined for the ATLAS experiment, which will search for theoretical Higgs Boson particles. But when tool co-developer and Lawrence Berkeley senior scientist Carl Haber tuned into a Public Radio report discussing the problem of how delicate or damaged media could be further degraded by contact with a stylus, he made a grand conceptual connection.

Haber and colleague Vitaliy Fadeyev went to Down Home Music in El Cerrito, Calif., and bought some 78rpm records. Using a high magnification Smart Scope optical camera at a vertical resolution of 100nm, the pair scanned the records 40 minutes for each second of audio. (Haber points out that a tool actually optimized for the task could likely scan a 10-inch 78rpm disc in five to 15 minutes.) With the groove maps in hand, they analyzed the images to reconstruct both the data and how the stylus originally passed along the groove.

"When we hear noise in an old recording as a 'click' or a 'pop,' that is due to the stylus striking a defect or some foreign material," notes Haber. "Present digital noise-reduction methods recognize that as an anomalous spike in the audio waveform and subtract it. In our approach, this would be treated as an anomaly in the image so it can be corrected on that basis before the 'virtual stylus' ever strikes it. If a recording is broken, cracked, or has missing pieces, it can't be played at all with a stylus. In our method, the separate pieces could be reassembled digitally."

Haber's group is now working with the U.S. Library of Congress to digitize early sound recordings in the national collection, including very rare Edison cylinders. ▲

Nano-Lightning Strikes Chip Heat

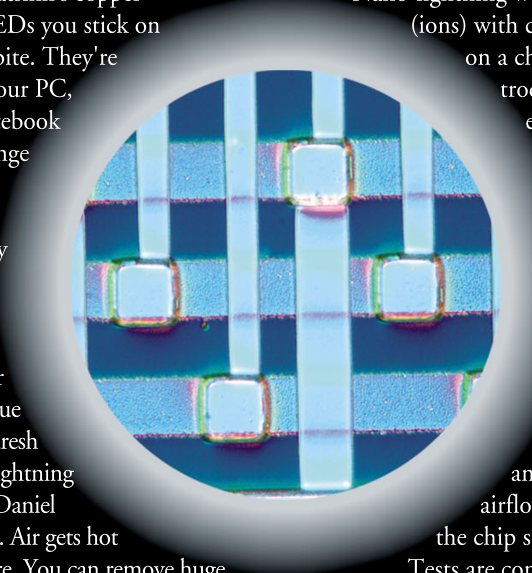

No matter how awesome a heatsink's copper sculpting is or how many LEDs you stick on it, the fact remains: Big heatsinks bite. They're heavy, chew up real estate inside your PC, and are a significant limiter on notebook form factors. Moreover, the challenge of aircooling only worsens as chips become faster and hotter.

Researchers at Purdue University think they may have a breath of fresh air for the heatsink industry in the form of "nano-lightning."

"It is much easier to pump hot air out than to pump liquid," says Purdue mechanical engineering professor Suresh Garimella, who invented the nano-lightning technology with associate professor Daniel Schlitz. "But aircooling has its limits. Air gets hot very easily; then it can't cool anymore. You can remove huge amounts of heat with liquid—a hundred, maybe a thousand times more. So the key attribute . . . is that it sticks with aircooling still, and yet, we're able to get liquid cooling-type rates of heat removal."

Nano-lightning works by creating electrically charged atoms (ions) with closely packed carbon nanotubes arrayed on a chip's surface. The nanotubes act as electrodes, with negatively charged electrodes emitting electrons toward the positively charged nanotubes. On a larger scale, this process results in atmospheric lightning. In this case, though, while there is no spark, the researchers are harnessing the "corona wind" phenomenon, or the directional movement of ions across an electrical field. By switching the electrode states and "pulling" ions where they want, researchers can control the corona wind across the chip surface and direct it through microchannels. The airflow pulls heat along with it until it leaves the chip surface.

Tests are continuing throughout the summer to determine actual rates of heat removal, but Garimella believes diamond-thin films can replace the nanotubes and the technology can be manufactured with conventional methods. ▲

Talk To The Card

The engineers behind Beepcard (www.beepcard.com) are putting the finishing touches on a technology that may soon have you protecting yourself from identity theft and fraud by speaking to your credit card.

The story starts with Beepcard's ComDot, a credit card with an integrated speaker and activation button. A user presses the button, and the speaker emits an electronic "squawk" much like a fax signal. A PC microphone picks up the squawk and transmits the signal to a server for authentication so that the transaction can proceed. Essentially, the ComDot is useful for online transactions in which the seller wants proof that the buyer actually owns the credit card being used. Additionally, the squawk code changes according to a prearranged sequence so that recorded squawks can't be used for fraud. Beepcard's next-generation VoiceCard technology will add a tiny microphone and a voice recognition chip into the equation.

"In order to do a login, you'll have to tell your credit card your password," says Beepard CEO Alan Sege. "Your credit card analyzes whether it was you and whether it was your password. If both of those things are true, then it will send its one-time number as a sound. Compliant sites have an ActiveX control running that listens for the card."

While resident data on the card is protected with 128-bit encryption, the real challenge for designers was the battery. Sege notes that a card's battery life is targeted at two to three years, or 8,000 operations, which is generally longer than it takes a card's magnetic strip to wear out.

Sege says the VoiceCard is all but done. All that's left is to finish the programming that will let Beepcard squeeze the equipment into a credit card-thin package. From there, the sky's the limit for applications.

"The voice product will do a lot of other things beside security," Sege says. "It has a speaker-independent vocabulary of up to 200 words. We can put programs in there like calculating a tip. Or you can pick up your card and say, 'Tell me a joke,' and it'll say, 'Your bank account.'" ▲

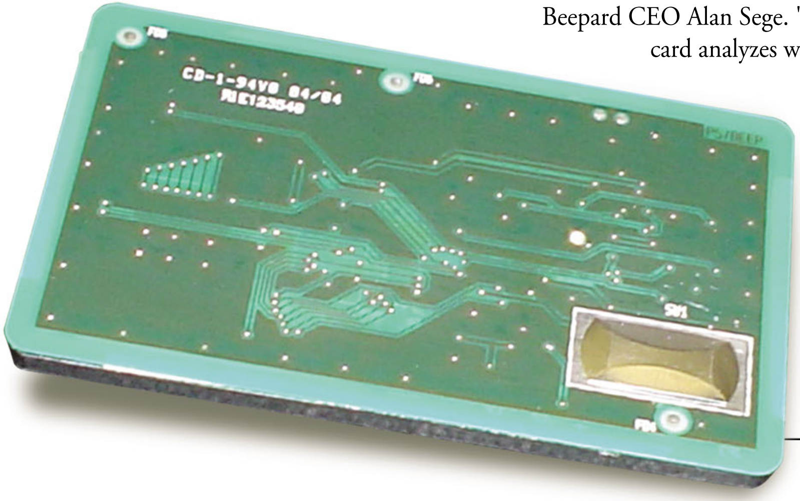

With a Beepcard VoiceCard in hand, your credit account will be protected by your voiceprint, as well as your PIN and 128-bit encryption, making fraud virtually impossible.

Q&A With Howard Rheingold

Prolific author, editor, philosopher, and critic, Howard Rheingold has been called The First Citizen of the Internet. He decided to ditch his typewriter in the early days of covering Xerox PARC and found a new direction in life upon connecting a telephone to his PC. Perhaps Rheingold might best be described as one of the world's first digital anthropologists, studying how we create the virtual world and what it does to us in return.

Q Tell us about your latest book, "Smart Mobs."

RHEINGOLD: Smart mobs are groups of people who use the Internet and mobile communications to coordinate and organize collective action. That's everything from teenagers who all show up at the mall at the same time because they've been exchanging text messages to citizens who demonstrate and bring down governments, as they did in the Philippines, or get out the vote and tip elections, as they did in Korea. In a larger sense, this is really about the merger of the mobile telephone, the personal computer, and the Internet. When you combine media, sometimes you get something new that has its own properties. I think that the unique characteristic of this new medium that's emerging from the merger of these prior technologies is the ability to organize collective action on new scales in new places with people who weren't able to organize before.

Q You probably saw the news story a while back about grade schoolers turning in papers written in SMS shorthand rather than regular English because that's how they think words are actually spelled. What's your take on that?

RHEINGOLD: Language is something that grows from the way people use it. At least they're writing things and not just being passive recipients of media.

Q What is the difference between a technology user and a technology consumer?

RHEINGOLD: The industry and the medium around the personal computer weren't created by the existing computer industry. It was created by 19-year-old dropouts like Steve Jobs and Bill Gates and a lot of other users who found ways to create uses for the technology, from the spreadsheet on. Those were really invented not by companies but by individual users. What is the Web but millions of people putting up Web sites? Those were not created by some central authority or some company but by a lot of people. Those are users. They are actively shaping the technology.

Now, in the old days when there was one telephone company and three television networks, people were passive consumers of what someone else created for them. The question about new technology is whether the attempt by existing companies to protect their business models by legislating or even changing the way the technology works is going to limit innovation to only people who work for those companies. The arguments about digital rights management and trusted computing, about the extension of copyright law, about the way wireless spectrum is regulated, those are essentially about innovation. In the future, is a 19-year-old dropout going to be able to use handheld technology to create a new industry or is he going to have to work for an existing company in order to be able to innovate?

Q You've written that, in contradiction to the 1996 Telecommunications Act, the cable companies and telcos have successfully petitioned the FCC to no longer have to carry traffic from competing networks. This obviously opens the door for the likes of Comcast, Time Warner, and Verizon to censor content from their competitors; does it also enable them to censor according to ideology?

RHEINGOLD: Well, we've seen that when the rules were changed about the ownership of radio stations, instead of owning a hundred stations, you could become a Clear Channel and own thousands. Clear Channel already has censored the content that it allows to be broadcast according to their political views. I don't see why that wouldn't apply to anyone else who has that power. The point is not what the ideology of the censoring party is, but whether they have the power to do it.

Q Do you think that ISPs will take steps like that in America—or have they already?

RHEINGOLD: Well, they can be forced to. The recording industry is more and more trying to force ISPs to police what people do with their networks. What if the FBI had made the telephone company responsible for anyone who had committed a crime using a telephone? Where would that technology have gone?

Subscribers can see www.cpumag.com/cpuaug04/rheingold for more content.

William Van Winkle began writing for computer magazines in 1996. He was first published in 1990, the same year he took his first job in computers. He and his family live outside Portland, Ore.

Quietly Setting A Whole New Standard.

Experience the unmatched combination of noise reduction and audio performance of our QuietComfort® 2 Acoustic Noise Cancelling® headphones.

Unwanted noise is everywhere. The engine roar inside an airplane cabin. The blaring sounds of city streets. The annoying din of the office. Bose® QuietComfort® 2 Acoustic Noise Cancelling headphones let you escape it all. Put them on and slip into a safe haven – where you can relax and enjoy peace and solitude or listen to your favorite music.

Proven Noise Reduction Technology.

Our original noise-reducing headphones were designed primarily for airplane travelers. But owners soon started telling us how much they enjoyed them in other noisy places. So we designed our QuietComfort® 2 headphones around the same patented technology that electronically identifies and dramatically reduces noise, while faithfully preserving the music, movie dialogue or tranquility you desire. The *Philadelphia Daily News* says that "Even in the noisiest environment, wearing these phones creates a calming, quiet zone for easy listening or just snoozing."

Our Best Sound Ever.

QuietComfort® 2 headphones deliver the best audio performance we've ever developed for headphones. Highs are impressively crisp and clean. Lows are deep and enveloping. Vocals are reproduced with lifelike authenticity. When audio critic Wayne Thompson heard our QuietComfort® 2 headphones, he reported, "Bose engineers have made major improvements." Whether you're looking for noise reduction or high-performance headphones for music, we think you'll agree – QuietComfort® 2 headphones simply make your music sound better.

Enhanced Comfort and Portability.

Respected columnist Rich Warren says, "The QuietComfort 2 lives up to its name, enveloping you in blissful sound in the utmost comfort. It's easy to forget they're on your head." To enjoy peace and quiet, simply turn them on. To add Bose quality sound, attach the included audio cord and connect them to a home stereo, laptop computer, portable CD/DVD/MP3 player or in-flight audio system. When you're done, their convenient fold-flat design allows you to quickly tuck your headphones away in a bag or briefcase, so you can take them almost anywhere.

Experience Them for Yourself.

We don't expect you to take our word for how dramatically our headphones reduce noise, how clean and full they sound or how comfortably they fit. So we invite you to try them for 30 days, risk free. If you can live without them, return them for a full refund. It's that simple. When you call, be sure to ask about our **financing options.**

QuietComfort® 2 headphones – the new standard from Bose, the most respected name in sound.

(Mr./Ms./Mrs.)
Name *(Please Print)*
Address
City/State/Zip
Day Phone Evening Phone
E-mail (Optional)
Mail to: Bose Corporation, Dept. DMG-Q3050, The Mountain, Framingham, MA 01701-9168

CALL 1-800-404-2073
ext. Q3050

For information on all our products: **www.bose.com/q3050**

BOSE®
Better sound through research®

©2004 Bose Corporation. Patent rights issued and/or pending. Financing not to be combined with any other offer or applied to previous purchases, and subject to change without notice. Risk free refers to 30-day trial only. Delivery is subject to product availability. Quotes are reprinted with permission: Jonathan Takiff, *Philadelphia Daily News*, 5/29/03; Wayne Thompson, *Oregonian*, 4/18/03; Rich Warren, *News-Gazette*, 5/19/03.

Super wireless performance–
it runs in our family.

Faster access to multi-player gaming
Video streaming without jitters
Effortless wireless connection – gets anyone up and surfing quickly
Superior wireless range

Total 802.11 Super G™ ATHEROS

Super G™ Technology – 10X Faster

108 Mbps 2.4 GHz 802.11 g

WGT624
108 Mbps
Firewall Router

WG511T
108 Mbps
PC Card

WAG511
108 Mbps
Dual Band
PC Card

WAG311
108 Mbps
Dual Band
PCI Adapter

Introduce your family to our NETGEAR® 802.11g family that offers exceptional range, compatibility, security and super speed.

Now you can choose 802.11g technology with confidence, thanks to NETGEAR. Recipient of two CNET awards for exceptional performance and ease of use, our 802.11g family runs circles around other companies' wireless networking products. It's fully compliant with the recently ratified 802.11g standard and compatible with all 802.11b and 802.11g products.

Our NETGEAR 802.11g family. The perfect addition to *your* family. For details, stop by for a visit at **www.netgear.com/go/g**.

54 Mbps 2.4 GHz 802.11 g

cnet Editors' Choice APR 2003
WG602
54 Mbps
Access Point

cnet Editors' Choice APR 2003
WG511
54 Mbps
PC Card

WG311
54 Mbps
PCI Adapter

WGR614
54 Mbps
Router

WGE101
54 Mbps
Ethernet Bridge

WG121
54 Mbps
USB Adapter

Everybody's connecting.™

NETGEAR®

©2003 NETGEAR, Inc. NETGEAR is a registered trademark and "Everybody's connecting." is a trademark of NETGEAR, Inc. All rights reserved. Other brands and products are trademarks of their respective holder/s. E134 Smart Computing